

Daetz, Stine Louise; Hvid, Anna Kirstine; Martinello, Alessandro; Matin, Rastin

Working Paper

Seeing through the spin: The effect of news sentiment on firms' stock market performance

Danmarks Nationalbank Working Papers, No. 141

Provided in Cooperation with:

Danmarks Nationalbank, Copenhagen

Suggested Citation: Daetz, Stine Louise; Hvid, Anna Kirstine; Martinello, Alessandro; Matin, Rastin (2019) : Seeing through the spin: The effect of news sentiment on firms' stock market performance, Danmarks Nationalbank Working Papers, No. 141, Danmarks Nationalbank, Copenhagen

This Version is available at:

<https://hdl.handle.net/10419/227859>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DANMARKS NATIONALBANK

4 OCTOBER 2019 — NO. 141

Seeing Through the Spin: The Effect of News Senti- ment on Firms' Stock Market Performance

Stine Louise Daetz
sld@nationalbanken.dk
UNIVERSITY OF COPENHAGEN

Anna Kirstine Hvid
akhv@nationalbanken.dk
UNIVERSITY OF COPENHAGEN

Alessandro Martinello
alem@nationalbanken.dk
DANMARKS NATIONALBANK

Rastin Matin
rma@nationalbanken.dk
DANMARKS NATIONALBANK

The Working Papers of Danmarks Nationalbank describe research and development, often still ongoing, as a contribution to the professional debate.

The viewpoints and conclusions stated are the responsibility of the individual contributors, and do not necessarily reflect the views of Danmarks Nationalbank.

Seeing Through the Spin: The Effect of News Sentiment on Firms' Stock Market Performance

Abstract

The sentiment of news predicts the short-term stock market performance of individual companies. We find that this association is solely due to the idiosyncratic informational content of an article. We transparently quantify the association between news sentiment and stock market performance of S&P 500 companies, using articles written by Reuters between 2000 and 2018. First, we isolate the effect of sentiment independently of idiosyncratic informational content by exploiting a topic-based shift-share instrument. Second, we show that exogenous variation in article sentiment isolated through our topic-based shift-share instrument, while strongly related to article sentiment, is unrelated to abnormal returns in the stock market.

Resume

Nyheders sentiment forudsiger aktiemarkedet på kort sigt for individuelle virksomheder. Vi viser, at denne sammenhæng udelukkende skyldes en artikels idiosynkratiske information. Vi undersøger sammenhængen mellem nyhedssentiment og aktiemarked for S&P 500 virksomheder ved brug af Reuters nyhedsartikler i perioden 2000 til 2018. Først isolerer vi effekten af sentiment uafhængig af idiosynkratisk information ved brug af et emne-baseret shift-share instrument. Dernæst viser vi ved brug af dette shift-share instrument, at den eksogene variation i artikel-sentiment ikke er relateret til anormale afkast, på trods af at det er kraftigt relateret til artikel-sentimentet.

Key words

Financial markets, Financial sector, Forecasting,

Statistical method

JEL classification

C55, D53, G15, G17

Acknowledgements

The authors are grateful to seminar participants at Danmarks Nationalbank for helpful comments.

The authors alone are responsible for any remaining errors.

SEEING THROUGH THE SPIN: THE EFFECT OF NEWS SENTIMENT ON FIRMS' STOCK MARKET PERFORMANCE*

S. L. Daetz, A. K. Hvid, A. Martinello and R. Matin

October 3, 2019
PRELIMINARY DRAFT

Abstract

The sentiment of news predicts the short-term stock market performance of individual companies. We find that this association is solely due to the idiosyncratic informational content of an article. We transparently quantify the association between news sentiment and stock market performance of S&P 500 companies, using articles written by Reuters between 2000 and 2018. First, we isolate the effect of sentiment independently of idiosyncratic informational content by exploiting a topic-based shift-share instrument. Second, we show that exogenous variation in article sentiment isolated through our topic-based shift-share instrument, while strongly related to article sentiment, is unrelated to abnormal returns in the stock market.

*Disclaimer: The views expressed in this paper are those of the authors and do not necessarily reflect the position of Danmarks Nationalbank.

The sentiment of stories reported by the media is strongly associated with stock market reactions. In the aggregate, Tetlock (2007) shows that negative words in finance columns in the *Wall Street Journal* 1984 to 1999 predict average stock returns on the Dow Jones Industrial Average, and Tetlock et al. (2008) shows that negative sentiment of articles mentioning specific firms predicts those firms' stock returns in the short run between 1980 and 2004.

Despite the robustness of these findings, the mechanism linking media sentiment and stock market performance is unclear. From a theoretical perspective, fundamental information drives stock market performance, and media sentiment is simply a proxy of that otherwise unobservable quantity. However, media sentiment may not only capture underlying fundamental information. Studies such as DellaVigna and Kaplan (2007) suggest that media sentiment directly and causally affects individual behavior. Moreover, media sentiment may be related to systematic fluctuations in investor sentiment (e.g. during financial crises) and behavioral shocks that are unrelated to stock fundamentals. As a consequence, such a direct causal link between media sentiment and stock market performance could amplify systematic fluctuations in the business cycle, inflate booms and busts of sentiment spirals, and undermine the stability of financial markets.

In this paper we test for the presence of a direct causal effect of news sentiment, independently of news' firm-related informational content, on firm-specific stock market performance, and we find that sentiment has no causal impact on its own. We study all articles about S&P 500 companies published between 2000 and 2018 by Reuters News, a global leader in delivering news to financial actors.¹ We quantify the sentiment of each of those articles, and link it to the stock market performance of each related company.

We proceed in two steps. First, we transparently and directly dissect the association between the stock market performance and article sentiment. Our approach, essentially a difference-in-differences comparison, focuses on robustly estimating the magnitude of the association rather than its predictive power. We compare the effect of positive and negative news on abnormal returns on the stock of a specific company ten trading days before and after the time of news publication.

We find that the association between news sentiment and firm-specific stock market performance is strong, both economically and statistically, but is decreasing over time. The largest (imprecise) effects occur before 2003, and 2017 marks the lowest average effect in our sample. Overall, a standard deviation increase in sentiment negativity is associated with a 25 basis point drop in abnormal returns on the day of news publication. Consistent with the results of Tetlock (2011), the drop is partially anticipated by the market the day before the release of the news. Our results are robust to alternative specifications, and are not due to general equilibrium effects. Compared to the periods studied in previous studies, this paper focuses on a time when news was easily accessible by all via internet. We therefore contribute to the literature by replicating earlier findings in a different global environment.

Second, we isolate exogenous variation in article sentiment exploiting a topic-driven shift-share instrument, and through this instrument test the existence of a causal relationship between media sentiment and firms' stock market returns. Once we isolate variation in media sentiment independent from idiosyncratic informational content about a specific firm, we find no relationship between media sentiment and stock market performance.

For identification, we exploit a shift-share instrument constructed on the basis of a topic decomposition of our article corpus. Shift share instruments are a popular identification strategy in

¹Reuters News reaches one billion individuals each day, and according to corporate sources, its associated trading platform Eikon has a 34% market share. Reuters is thus a massive player in financial information delivery, with the potential for general equilibrium disruption.

applied econometrics, with applications ranging from labor economics (Bartik, 1991; Jaeger et al., 2018) to economics of the household (Aizer, 2010) to macroeconomics (Nakamura and Steinsson, 2014). This type of identification strategy formalizes the identifying assumption that innovations over time across separate trends affect specific units differently, depending on the predetermined exposure of each unit to each trend.

For example, Foged and Peri (2016) exploits that flows of immigrants of different nationalities affected specific Danish municipalities according to the ethnic composition of their immigrant population, as immigrants tend to cluster in communities. Nakamura and Steinsson (2014) exploits that changes in federal spending for military procurement affected specific US regions according to predetermined military procurement build-ups. Autor et al. (2013) exploits that changes in imports of goods from China affected specific US regions according to their exposure to China imports, in turn determined by their industry composition. In most of these applications, the identification assumption is that the variable determining the exposure to a time-varying trend (e.g. industry composition) does not predict innovations to the outcome of interest (Goldsmith-Pinkham et al., 2018).

In our case, we exploit that changes in the aggregated sentiment of different news topics affect specific firms differently, and according to the predetermined exposure of each firm to a specific news topic. We identify news topics in our data by decomposing our article corpus through a Latent Dirichlet Allocation (LDA) model (Blei et al., 2003). For each topic, we compute the exposure of each firm to a topic share using articles predating 2005. We use these shares to decompose how each firm is affected by weekly innovations to the aggregated sentiment news topics (explicitly excluding contributions to sentiment of articles mentioning the firm of interest). Our identifying assumption is that news topic shares determined before 2005 do not predict changes in abnormal returns at the time of article publication between 2005 and 2018.

While our instrument is strongly related to article sentiment, it has no detectable association with changes in abnormal returns. Our results imply that for the past decade news sentiment has had no direct impact on stock market performance over and beyond its idiosyncratic informational content.

Our work is related to the literature on investor sentiment (e.g. Baker and Wurgler (2006) and Baker and Wurgler (2007)). In this literature, investor sentiment is defined as a belief about future cash flows and investment risks that is not justified by the facts at hand, and is typically measured using different kinds of proxies (e.g. a combination of closed-end fund discount, NYSE share turnover, number of first-day returns on IPOs, equity share in new issues and dividend premium). While the media sentiment that we compute is a linguistic measure that should not be confused with the broader notion captured by this definition of investor sentiment, by identifying exogenous variation in media sentiment we are able to precisely capture one facet of investor sentiment.

Finally, note that the causal effect we focus on in this paper is that of media sentiment, and not of media availability. There is strong evidence that availability of information affects both individual behavior (Gerber et al., 2009) and the stock market. For example, Engelberg and Parsons (2011) convincingly show that earning announcements, when and where exogenously not reported by local media, affect stock markets differently. Our paper adds to their contribution by focusing on article content, rather than availability, which was widespread in the period we study via the internet.

The remainder of the paper is structured as follows. Section 1 describes our data sources and how we process the unstructured text data into sentiment scores. Section 2 dissects the association between article sentiment and stock market performance, studying heterogeneity over time, sector, and week of earning announcement. Section 3 describes our identification strategy and presents our

causal results. Section 4 concludes.

1 Data

1.1 News Articles from Reuters

Our initial article corpus consists of more than 13 million articles from Thomson Reuters News Archive. The news data includes metadata describing date of publication, topic- and entity tags, and language.

We restrict the analysis to news articles written in English, and apply straightforward filters to remove entries that summarize different unrelated news or simply report tables of stock market returns. If there are subsequent corrections to an article, we use the first version of the article within a 12 hour period, and in case of additions to an article within a trading day, we use the article with the longest body text.²

We focus on articles about firms that have been and will be part of the S&P 500 index for at least ten days since the time of publication. Matching articles to specific S&P 500 companies is not straightforward. The entity tags identifying a company mentioned in an article consist of a RIC (Reuters Instrument Code), which can change over time. We match company names and RICs over time using a regex search, and then, based on those RICs, extract tagged articles that explicitly mention at least one of possibly multiple company names (e.g. Google or Alphabet) in the title. Our final article selection consists of 288814 articles published between 2000 and 2018 (on average 41 each day).

We focus on two measures of article sentiment. The simplest, and standard in the finance literature, is a dictionary-based approach developed by [Loughran and McDonald \(2011\)](#). This dictionary³ contains 354 positive words and 2355 negative words. The resulting sentiment measure according to this approach is a simple difference between the negative and positive word count in each article, normalized by the total number of words in the article. Therefore, the higher is the measure, the more pessimistic is an article. Compared to other dictionary approaches, this dictionary has the advantage of having been specifically compiled for analyzing finance text, and therefore includes words that might have a different meaning in a non-finance context. For example, the words “antitrust” and “concedes” have a negative meaning according to this dictionary, while they might be considered neutral or even positive in other contexts.

This first approach has the disadvantage of not being particularly nuanced. Dictionary approaches can’t understand lexical constructs as negations and degree modifiers. The constructs “not catastrophic” and “very catastrophic” would receive the same score of 0.5, due to the presence of the negative word “catastrophic”. We therefore exploit the VADER algorithm designed by [Hutto and Gilbert \(2014\)](#). We construct the VADER score as the difference between the negative and positive score that VADER assigns to a sentence. For example, the two constructs mentioned above would obtain a score of -0.724 (optimistic sentence) and 0.777 (pessimistic sentence).

The disadvantage of this second approach is however that, having been developed for the analysis of social media texts, it does not assign any particular meaning to words like “antitrust” and “concedes”, which are deemed as intrinsically neutral. Sentences like “CEO confesses committing

²When a news item breaks, Reuters typically publishes immediately a breaking news alert, often consisting of a single sentence. The body of the article is then added within a few minutes. In our corpus we observe both entries separately, but we use the second, updated version to compute the sentiment of the news.

³The dictionary is available online at <https://sraf.nd.edu/textual-analysis/resources/>

corruption”, which we would consider very bad news for any given company, are rated as positive in such a context. In our second sentiment measure we therefore combine the complementing strengths of the two approaches by taking the first principal component of the two scores. We normalize both measures such that they have a mean zero and standard deviation of one in our sample.

1.2 Stock Market Data

We restrict our sample to firms that at some point during the period January 2000 to March 2019⁴ were included in the S&P 500 index, and for which we observe at least five articles. Since the S&P index constitutes the 500 largest companies in terms of market capitalization, we ensure the presence of media attention on the firms, as well as the availability of firm-level stock price and balance sheet data. Our sample consists of 798 companies and for each firm we analyze the impact of negative news for periods during which the firm is covered in the S&P 500 index.⁵

We rely on data from Bloomberg to measure stock market performance. Specifically, we use daily end-of-day stock price information during our sample period for each firm in our main sample, and only use stock price information related to the New York Stock Exchange (NYSE). The NYSE is open for trading Monday through Friday from 9:30 am to 4:00 pm Eastern Time, with the exception of holidays declared by the exchange in advance.⁶

As we are interested in investigating the impact of negative news on firms’ stock market performance we use two measures to test the *abnormal* development of firms’ stock prices following a news announcement. First, we use a benchmark approach to measure changes in firms’ stock prices. As our baseline benchmark measure we use the difference between the relative change in firms’ stock prices and the relative change in the S&P 500 index price on a given date. To account for sector specific stock price movements, we also provide a sector specific benchmark measure that is determined as the difference between the relative change in firms’ stock prices and the relative change in the S&P 500 sector index with the sector being determined by the firms’ two-digit GICS code.

Second, we use a capital asset pricing model (CAPM) approach to measure changes in firms’ stock prices. Specifically, we obtain each stock’s expected return by determining the stock’s beta with the S&P 500 index over the period January 2000 to March 2019, and then measure the firm’s abnormal stock return using the difference between the actual and the expected stock return.⁷

We also investigate the impact of news on stock volatility and trading volume. Specifically, we proxy firms’ intraday stock volatility by taking the difference between the highest and the lowest stock price, and dividing it by the closing price. As the number of news articles correlates with the size of firms, we also obtain baseline balance sheet data from Bloomberg and use firm size, measured as the firm’s total assets given in natural logarithm, as a control variable.

Since we use news articles from Reuters, which are timed in accordance to central time standard, we combine the data from Reuters and Bloomberg by allocating news articles after 10 am CET to the stock price information from the following day. The remaining news articles are allocated to the

⁴Note that when stock market data is combined with Reuters data, we only include data until 2018.

⁵For those firms that appear in the S&P 500 index more than once, we only investigate the impact of negative news for the longest of these data periods.

⁶The NYSE averages about 253 trading days per year.

⁷In robustness checks we also use time-varying betas, i.e., beta values that are only based upon observations of the stock and the S&P 500 index during the previous year and, thus, allow for changes in the firm’s stock over time. Our results are robust when we use the time-varying beta in the CAPM estimation.

Table 1
Summary Statistics

| Sector | N | Stock Return | | | Intraday Vol. | | | Δ Volume | | | Firm Size | | |
|------------------------|--------------|--------------|-----------|-----------|---------------|-----------|-----------|-----------------|-----------|------------|------------|------------|------------|
| | Firms (1) | Mn (2) | Md (3) | Sd (4) | Mn (5) | Md (6) | Sd (7) | Mn (8) | Md (9) | Sd (10) | Mn (11) | Md (12) | Sd (13) |
| Communication Services | 47 | 0.04 | 0.03 | 0.06 | 0.03 | 0.03 | 0.01 | 0.19 | 0.12 | 0.36 | 9.61 | 9.59 | 1.29 |
| Consumer Discretionary | 102 | 0.06 | 0.06 | 0.05 | 0.03 | 0.03 | 0.01 | 0.16 | 0.13 | 0.18 | 8.67 | 8.65 | 1.04 |
| Consumer Staples | 51 | 0.05 | 0.04 | 0.05 | 0.02 | 0.02 | 0.00 | 0.13 | 0.12 | 0.08 | 9.19 | 9.19 | 1.21 |
| Energy | 63 | 0.06 | 0.06 | 0.05 | 0.03 | 0.03 | 0.01 | 0.12 | 0.11 | 0.08 | 9.13 | 8.93 | 1.04 |
| Financials | 105 | 0.05 | 0.05 | 0.04 | 0.03 | 0.03 | 0.01 | 0.23 | 0.11 | 1.04 | 10.80 | 10.79 | 1.70 |
| Health Care | 97 | 0.08 | 0.07 | 0.06 | 0.03 | 0.03 | 0.01 | 0.23 | 0.14 | 0.66 | 8.53 | 8.27 | 1.32 |
| Industrials | 95 | 0.06 | 0.06 | 0.08 | 0.03 | 0.03 | 0.02 | 0.17 | 0.12 | 0.38 | 8.88 | 8.62 | 1.23 |
| Information Technology | 119 | 0.06 | 0.06 | 0.04 | 0.04 | 0.03 | 0.01 | 0.20 | 0.13 | 0.46 | 8.37 | 8.07 | 1.32 |
| Materials | 52 | 0.05 | 0.05 | 0.05 | 0.03 | 0.03 | 0.02 | 0.28 | 0.12 | 0.89 | 8.78 | 8.72 | 1.09 |
| Real Estate | 29 | 0.05 | 0.03 | 0.04 | 0.02 | 0.02 | 0.01 | 0.16 | 0.15 | 0.07 | 8.88 | 8.82 | 0.62 |
| Utilities | 38 | 0.04 | 0.04 | 0.04 | 0.02 | 0.02 | 0.01 | 0.15 | 0.11 | 0.17 | 9.96 | 9.95 | 0.68 |
| Total | 798 | 0.06 | 0.06 | 0.05 | 0.03 | 0.03 | 0.12 | 0.19 | 0.12 | 0.56 | 9.12 | 8.98 | 1.46 |

NOTE: The table reports the number of firms in each sector, as well as the mean, median and standard deviation of stock returns, intraday volatility, change in volume and firm size. Intraday volatility is given by the difference between the highest and the lowest stock price divided by its closing price. The change in volume is given by the percentage difference in trading volume on the NYSE relative to the day before. Firm size is measured as the natural logarithm of a firm's total assets.

same day. We round up any events occurring on either a weekend or a US-holiday to the following trading day.

Table 1 outlines an overview of our firm sample by sector, including some base summary statistics for our outcome variables. Of the 798 firms in our sample, the IT, Financial -and Consumer Discretionary sectors are the most represented, and financial firms are the largest on average.

2 Difference in Differences Approach

This section shows that news is strongly associated with the stock market performance of S&P 500 companies, replicating earlier results in an event-study framework. Confirming the results of Tetlock et al. (2008) and Boudoukh et al. (2013), we show that sentiment is crucial to understanding and interpreting the effect of news in the stock market. Consistent with Tetlock (2007), our results suggest that news sentiment is partially anticipated in the market, highlighting its endogeneity.

This section proceeds in three steps. First, we transparently describe the association between news sentiment and stock market performance using a Difference-in-Differences (DiD) strategy. Second, we explore sources of heterogeneity in this association. The effect is stronger at times of

earning announcements and decreasing with time, with the smallest association recorded in 2017. Third, we quantify the predictive power of news sentiment in the financial market in a portfolio setting, following Tetlock et al. (2008). Consistently with our heterogeneity analysis, a portfolio strategy using news sentiment to rebalance portfolios would have yielded in the past five years less than a third of the returns it would have yielded between 2001 and 2005.

Our approach aims at transparently characterizing the association between news occurrence, news sentiment, and the stock market performance of a firm subject to a news story. To this end, rather than employing standard tools of return predictability (Novy-Marx, 2014), we begin by comparing average stock market outcomes measured in the ten trading days before and after the publication of a negative news story with those measured around the publication of a positive news story for the same firm.

This comparison, akin to a difference-in-differences (DiD) approach centered around an event time (the time of the news release), allows us to transparently measure the association between news sentiment and stock market performance in terms of e.g. abnormal returns and changes in trading volume. Formally, we model an outcome $y_{i,f,t}$ related to article i , firm f , and trading day t as

$$y_{i,f,t} = \sum_{n=-10}^9 \beta_n + \sum_{n=-10}^9 \gamma_n \mathbf{1}[n] \cdot S_i + \Phi_i + \Lambda_{i,t} + \varepsilon_{i,f,t} \quad (1)$$

where n is trading days from the occurrence of the event, and is simply the difference in trading days between the observed trading day and trading day during (or immediately following) the date of publication of article i . The coefficients of interest are therefore the set γ_n .

$\Lambda_{i,t}$ is a vector of controls consisting of fixed effects for day of the week of the trading date t (Monday to Friday) and firm size. We include fixed effects at the event-spell level Φ_i (which nest firm and article fixed effects), thereby adjusting for time-invariant characteristics of firm f that might be systematically correlated with article sentiment. We allow for arbitrary autocorrelation of errors $\varepsilon_{i,f,t}$ within all observations pertaining to the same firm and related to an article published in the same week (about 125 thousand clusters).

In our main DiD specification, S_i identifies negative events in our sample, defined as those belonging to the top 33% of the distribution in terms of negative sentiment. We define our control group accordingly: firms in the bottom 33% of the distribution in terms of negative sentiment in our baseline specification.

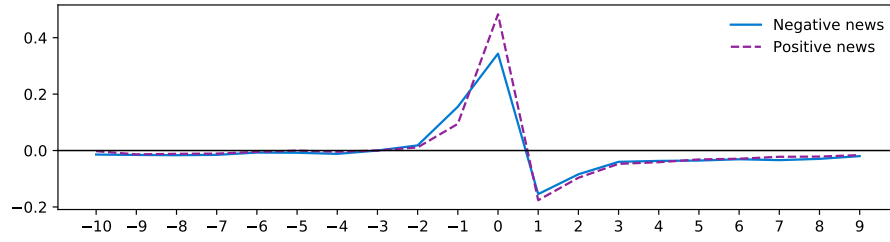
Figure 1 shows how trading volume, intraday volatility and abnormal returns evolve around the date of publication of a news story. Trading volume and volatility increase irrespective of the sentiment of the story, measured as the first principal component of our raw sentiment scores, indicating that articles carry information that the market has not already internalized.

However, abnormal returns drop when news is negative, and spike when news is positive, as sentiment captures some of the directional information contained in the article. The average difference in returns between positive and negative news on the day of publication is substantial, with a 0.52% difference in excess returns in a single trading day.⁸

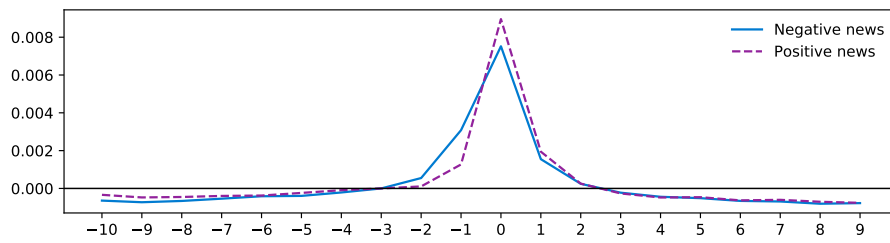
While transparent, our DiD approach does not allow us to precisely quantify the association between sentiment and stock market performance. To this purpose we estimate an alternative

⁸These qualitative results using a DiD approach are consistent across sentiment measures, and whether we compute abnormal returns as the residuals of a CAPM model or as the simple difference between stock returns and sectoral S&P 500. These results appear in Table 5 in the Appendix.

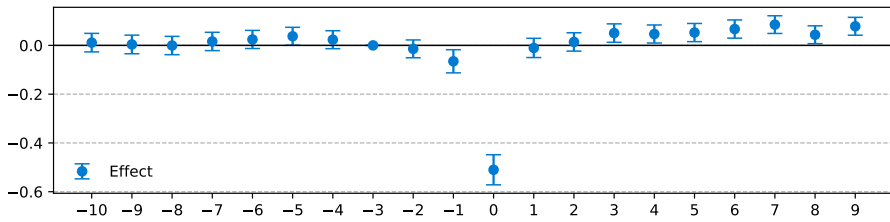
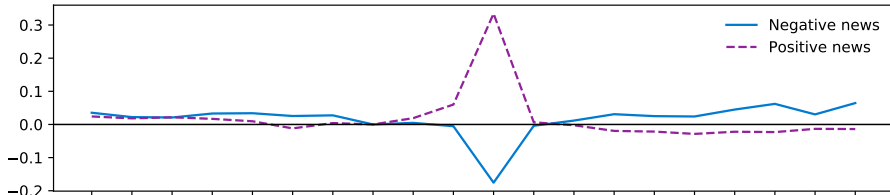
Figure 1
The effect of article publication on the stock market over time



(a) Change in trading volumes



(b) Intraday volatility



(c) Abnormal returns (CAPM)

NOTE: The top three panels plot the conditional average of outcomes between events characterized by the publication of positive and negative news, represented by coefficients β_n in equation (1). Positive and negative articles are defined as the bottom and top 33% of articles according to the principal component of the [Loughran and McDonald \(2011\)](#) and [Hutto and Gilbert \(2014\)](#) negativity scores. The bottom panel shows the estimated difference over time of abnormal returns, calculated as the residuals from a CAPM 3-factors model, corresponding to coefficients γ_n in equation (1).

linear specification of model (1) where S_i , rather than a group indicator, represents news sentiment standardized by its standard deviation. These results appear in Table 2.

The first column of the table reports the p-value of an F-test of joint equality to zero of all γ_n coefficients for $n < -4$ showing that days ahead of the event time the stock returns of firms affected by positive and negative upcoming news follow the same path. The remaining columns report estimated coefficients γ_n from model (1) for a sample of time-to-event n .

Table 2
Dynamic effects, linear specification

| n→ | P-val. | Change in outcome relative to $n = -3$ | | | | | | |
|---|--------|--|-------------------|-------------------|-------------------|------------------|------------------|------------------|
| | < -3 | -2 | -1 | 0 | 1 | 2 | 5 | 9 |
| <i>Panel A: Loughran/McDonald composite score</i> | | | | | | | | |
| Abnormal Returns (CAPM) | 0.937 | -0.009 (0.009) | -0.061 (0.010) | -0.228 (0.013) | -0.006 (0.009) | 0.015 (0.009) | 0.020 (0.009) | 0.027 (0.009) |
| Abnormal Returns (SD) | 0.927 | -0.017 (0.007) | -0.065 (0.009) | -0.231 (0.012) | -0.014 (0.008) | 0.006 (0.007) | 0.013 (0.007) | 0.019 (0.007) |
| - Own Returns | 0.936 | -0.009 (0.009) | -0.061 (0.010) | -0.228 (0.013) | -0.006 (0.009) | 0.015 (0.009) | 0.020 (0.009) | 0.027 (0.009) |
| - Sector Returns | 0.479 | 0.008 (0.005) | 0.004 (0.005) | 0.003 (0.005) | 0.008 (0.005) | 0.009 (0.005) | 0.007 (0.005) | 0.008 (0.005) |
| Book to market | 0.573 | 0.062 (0.025) | 0.183 (0.044) | 0.353 (0.082) | 0.402 (0.192) | 0.490 (0.254) | 0.496 (0.245) | 0.495 (0.249) |
| <i>Panel B: Principal component</i> | | | | | | | | |
| Abnormal Returns (CAPM) | 0.698 | -0.004 (0.008) | -0.055 (0.010) | -0.252 (0.013) | -0.005 (0.009) | 0.011 (0.009) | 0.022 (0.009) | 0.028 (0.008) |
| Abnormal Returns (SD) | 0.900 | -0.009 (0.007) | -0.056 (0.008) | -0.251 (0.012) | -0.010 (0.007) | 0.009 (0.007) | 0.014 (0.007) | 0.023 (0.007) |
| - Own Returns | 0.696 | -0.004 (0.008) | -0.055 (0.010) | -0.252 (0.013) | -0.005 (0.009) | 0.011 (0.009) | 0.022 (0.009) | 0.028 (0.008) |
| - Sector Returns | 0.803 | 0.004 (0.005) | 0.001 (0.005) | -0.001 (0.005) | 0.005 (0.005) | 0.002 (0.005) | 0.008 (0.005) | 0.005 (0.005) |
| Book to market | 0.184 | 0.064 (0.024) | 0.172 (0.048) | 0.337 (0.088) | 0.340 (0.168) | 0.415 (0.214) | 0.439 (0.213) | 0.400 (0.215) |

NOTE: The table reports the estimated $\hat{\gamma}_n$ coefficients around the time of news publications for a series of stock market outcomes. Standard errors clustered by firm-by-week of publication level are reported in parentheses. The first column reports the p-value of an F-test testing joint equality to zero of all γ_n coefficients for $n < -3$. We can't reject the hypothesis of parallel pre-event trends for any of the outcomes. The first panel uses as a sentiment measure the Loughran/McDonald score; the second, the principal component of the Loughran/McDonald and VADER scores. Both scores are standardized.

The table shows that a standard deviation increase in the negativity of a news item is associated with a drop in abnormal returns on the day of news publication of between 0.2% and 0.26%. There is no difference in the evolution of returns three or more days before the news publication. However, our data shows that even the day before the publication of a news item the market begins to react, consistent with the finding of [Tetlock \(2010\)](#). This early reaction stresses how the association between news sentiment and stock market returns is at least partly due to existing information in

the market, of which news sentiment is a (sometimes delayed) proxy.

The abnormal results do not persist in the days after the publication of the news. If anything, abnormal returns marginally increase for stocks in the weeks following negative news, likely due to progressive regression to the mean. Nonetheless, the effect of stock prices is persistent. Following negative news, the book to market ratio increases, signaling a lower market evaluation, and the effect is still significant after nine trading days since the occurrence of the negative news.

These results are robust to different specifications. The table shows that if we compute abnormal returns as the simple difference between a stock’s own return and those of the corresponding sector rather than as the product of a CAPM model, our results remain the same. Similarly, our results do not change substantially according to the measure of news sentiment we use. Both negativity computed exclusively via the dictionary developed by [Loughran and McDonald \(2011\)](#) (LM composite score) and the principal component of the LM score and the score computed via VADER ([Hutto and Gilbert, 2014](#)) deliver similar results. However, as expected, the principal component of the two measures is a slightly stronger predictor of stock market returns. In the remainder of the paper we focus on this composite measure for our analysis.

Our results are not due to general equilibrium effects. One could imagine that due to negative news about a specific company, the market would adjust by moving capital from that firm to others in the same sector, thereby amplifying the gap in returns. We do not find any evidence of such a mechanism. When negative news affects a specific company, the average returns of S&P 500 firms belonging to the same sector do not change. All of the impact on abnormal returns are due to changes in the affected firm’s stock returns.

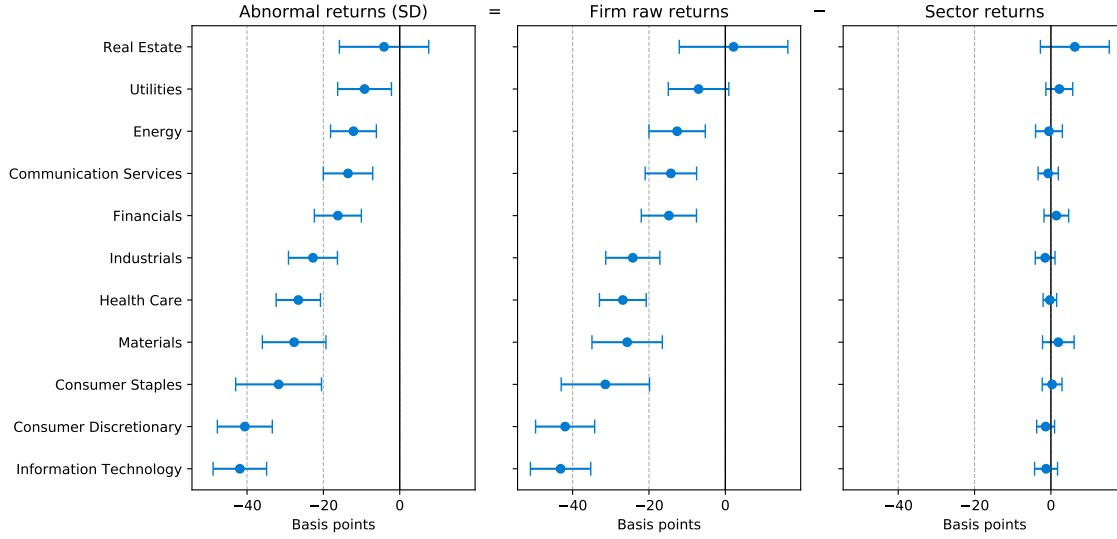
Similarly, while the effect of news on abnormal returns varies across sectors, we do not find any detectable rebalancing effect even within a single sector. The effect of news sentiment is most pronounced for the IT and consumer goods sectors, and smallest for real estate companies. Real estate is however also the sector for which we observe the least number of articles (4153). The second smallest sector in terms of news coverage is the energy sector (10776 articles), while the financial sector is the most covered (42247 articles).

We also verify the relationship between the effect of news sentiment and the timing of earning announcements by a firm. Earning announcements are particularly relevant events in the financial year of a company, and [Engelberg and Parsons \(2011\)](#) explicitly focus on news covering earning announcements. We find that the effect of news released in the same period of firm’s earning announcements is 50% stronger than that of news released outside of earning announcement windows.⁹

Finally, we show that the association between news sentiment and abnormal returns in the stock market is decreasing over time, and was at its lowest in 2017. Figure 3 shows that the effect of news released in the past few years is in general less than half with respect to news released before 2004. Such a decrease is consistent with multiple hypotheses. For example, due to automated trading shocks are absorbed faster by the market in recent years, and thereby we can’t observe fluctuations with day-to-day returns. Alternatively, as [McLean and Pontiff \(2016\)](#) suggests, by being aware of possible abnormal returns available by trading according to news sentiment, markets can incorporate and thus destroy return predictability of observable factors. Indeed, the next section shows that a trading strategy based on news sentiment would have been profitable up to recent years.

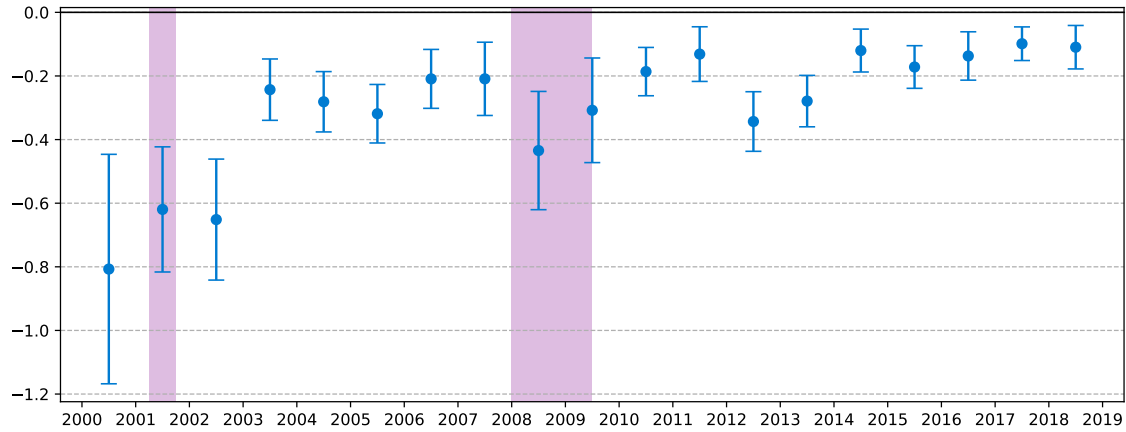
⁹These results appear in Table 6 in the Appendix.

Figure 2
Heterogeneity of the association between news sentiment and abnormal returns by firm sector



NOTE: The figure plots estimated $\hat{\gamma}_0$ coefficients according to the specification in model (1), using the standardized continuous principal component of the Laughran/McDonald and VADER scores as a measure of negativity, in different subsamples. Standard errors are clustered at the firm-by-week of article publication level.

Figure 3
Heterogeneity of the association between news sentiment and abnormal returns by calendar year



NOTE: The figure plots estimated $\hat{\gamma}_0$ coefficients according to the specification in model (1), using the standardized continuous principal component of the Laughran/McDonald and VADER scores as a measure of negativity, in different years. Standard errors are clustered at the firm-by-week of article publication level.

Table 3
Annual raw returns from news based trading strategy

| 1996-2000 | 2001-2005 | 2006-2010 | 2011-2015 | 2016-2019 | Overall |
|-----------|-----------|-----------|-----------|-----------|---------|
| 22.14% | 30.55% | 22.42% | 16.31% | 9.89% | 21.71% |

NOTE: The table reports the cumulative raw returns of a long-short strategy, where the long and short portfolio is held for one full trading day and rebalanced at the end of the next trading day.

2.1 Using Media News to Predict Stock Returns

The previous section demonstrates the negative relationship between negative media news and firms' stock market returns. In order to quantify the predictive power of negative news for stock returns, we turn to a predictive regression framework such as that used by Tetlock et al. (2008). The underlining hypothesis is that investors do not immediately respond in full to negative news. Accordingly, one would expect a positive return from a trading strategy that exactly exploits the news based information.

Given our results, a fairly simple news based trading strategy could imply positive risk-adjusted profits. Given that investors may not respond immediately and, in particular, underreact to negative news in the short run, one may benefit from going short in stocks with negative news over a one-day horizon.

We test the prediction of positive stock returns from a news based trading strategy in a portfolio setting. Using end-of-day stock price observations, we form two equal-weighted portfolios based on the news content provided about each of the firms in our sample. Specifically, we use the same definition for negative and positive news as before and include all firms with negative news, i.e., at least one negative word in the news text, on the prior trading into the short portfolio, and all firms with positive news, i.e., no negative word in the news text, on the prior trading into the long portfolio. Then we hold the short and long portfolio for one full trading day, and rebalance at the end of the next trading day.¹⁰ We exclude firm-days where we do not have news information about the firm, and likewise also exclude potential days when we do not have any firm in either the short or long portfolio.

Table 3 provides an overview of the cumulative raw returns of this long-short strategy. Using the full trading period from January 2000 until March 2019, we find that the raw return (ignoring trading costs) is about 21.7% per year¹¹. As expected, we also find that the return has decreased over time to about only 9% in most recent years.

3 Identification of Causal Effects through a Shift-Share Instrument

In a frictionless financial market, the association between news sentiment and stock market performance would solely be due to the underlying information contained in the news article. News sentiment would then simply be a proxy for this otherwise unobserved information. However, if news

¹⁰For each firm in the short portfolio, we sell one share at the stock price at time t , and then buy back one share at the the stock price at time $t+1$. For each firm in the long portfolio, we do the opposite.

¹¹In robustness checks, we run the same exercise but assume positive transaction costs and find that for small round-trip costs, there is still a positive raw return. Specifically, when assuming 1, 2, 5, or 10 bps of round-trip costs, the raw return drops to 18.69, 15.74, 7.31, and -5.4% respectively.

sentiment affects investor sentiment, and the latter affects behavior in the financial market (Baker and Wurgler, 2007), then news negativity might amplify negative market fluctuations through a negative feedback circle. This section tests for the presence of this direct causal link between the stock market and news sentiment, independently from idiosyncratic information about a specific company.

To test for such a link, we construct a shift-share instrument to isolate exogenous variation in the sentiment of news driven by topic-specific variation in sentiment. We thus exploit information that, while affecting the negativity of a news story, does not differentially affect stock market returns around the time of news publication except through the news sentiment.

Shift-share instruments are a popular tool in applied microeconometrics at least since Bartik (1991). Most applications of this strategy are in the field of labor economics Autor et al. (2013); Jaeger et al. (2018), but recent applications range from economics of the household (Aizer, 2010) to macroeconomics (Nakamura and Steinsson, 2014). The identification provided by such shift-share instruments has been recently formalized by Goldsmith-Pinkham et al. (2018) and Borusyak et al. (2018), to whom we refer for a more formal description of this strategy.

A shift-share instrument isolates variation in the variable of interest by averaging exogenous shock with shock exposure weights. Traditionally, this instrument has been used to isolate changes in wages driven by exogenous labor-demand factors by averaging national-level shocks in industry-specific wages with regional-specific fixed industry shares, so that a region initially more exposed to an industry experiencing an economic downturn nation-wide will experience lower wage growth than a region initially more exposed to growing industries.

Alternatively, Autor et al. (2013) weight changes in US imports from China by industry within a region using regional industry-specific weights computed before the entrance of China in the WTO. Foged and Peri (2016) weight changes in Danish inflows of refugees by country of origin within a municipality using municipal nationality-specific weights (as immigrants tend to cluster by nationality). Nakamura and Steinsson (2014) weight changes in US military expenditure within a region with weights defined by beginning-of-period regional military build-ups exposure.

In our setup, we weight aggregate innovations in news topic sentiment over time using firm-specific topic exposure. Imagine two firms, A and B, where A is more exposed to news about trade than firm B. If in a specific week articles talking primarily about trade become relatively more negative, a shift-share instrument would predict that articles about firm A would be relatively more negative than article about firm B. To avoid reflection, for a specific firm we compute aggregate trends explicitly excluding articles mentioning that specific firm. The topic-driven, firm-weighted information is known in the market. With respect to the days before the publication of a story about firm A, we then assume that this information would not differentially affect stock market returns of firm A on the day of article publication, except by affecting the sentiment of the story in a way that is unrelated to information specific to firm A. In other words, we isolate a source of variation in the sentiment of news about firm A that is not dependent on idiosyncratic information about firm A.

Formally, for each firm f we construct the firm-specific average topic share of all articles mentioning that firm up to 2005-01-01. We denote this measure $z_{f,\tau,2005}$ and by construction we have that $\sum_{\tau} z_{f,\tau,2005} = 1$. We construct the shift-share instrument for news sentiment as

$$\tilde{S}_{w(i),f} = \sum_{\tau} z_{f,\tau,2005} \check{S}_{w(i),\tau,-f} \quad (2)$$

where $\check{S}_{w(i),\tau,-f}$ is the change with respect to 2005 in the aggregate measure of sentiment of topic τ

(excluding articles mentioning firm f to avoid reflection problems) for each observed calendar week $w(i)$ after 2005-01-01 in which an article was published. As a consequence, our instrument varies at the calendar week by firm level. We standardize $\tilde{S}_{w(i),f}$ such that it has a standard deviation of 1.

In order to compute our shift-share instrument, we must first identify the underlying topics in our news articles. We extract these topics by using a topic model, which is a statistical method for obtaining the latent topics that occur in a collection of documents. Specifically, we employ the unsupervised topic model known as Latent Dirichlet Allocation (LDA). Originally introduced by Blei et al. (2003), it has since grown to be one of the most commonly employed topic models (Zhao et al., 2015). The LDA procedure combines words into topics (i.e., word distributions), while simultaneously decomposing articles into linear combinations of topics, thereby condensing an immense data set. The procedure starts by standard preprocessing of the M documents as outlined in Appendix B.

As only the occurrence of words in a document is important for the LDA procedure, and not the exact locations, we construct a vocabulary \mathcal{V} based on the unique words in the documents. The underlying assumed data generating process is that a document is formed by drawing words from a collection of topics, each of them in turn representing a different distribution over a set of words.

Formally, we can represent a document as a distribution θ_m over all possible topics K , where a topic k is a multinomial distribution ϕ_k over all words in \mathcal{V} . The structural assumption is that θ_m and ϕ_k have conjugate Dirichlet distributions with parameters α and β . We set these parameters to the commonly used values of $\alpha = 1/K$ and $\beta = 0.1$. The entire news corpus consisting of M documents is then constructed with the following procedure:

1. For each document $m \in 1, \dots, M$ draw a multinomial distribution over the K topics from a Dirichlet distribution with concentration parameter α :

$$\theta_m \sim \text{Dirichlet}(\alpha)$$

2. For each topic $k \in 1, \dots, K$ draw a word distribution from a Dirichlet distribution with concentration parameter β :

$$\phi_k \sim \text{Dirichlet}(\beta)$$

3. For all corpus word positions i, j , where i indicates the i th document and j indicates the j th position of the word in document i :

- (a) Draw the topic the word originates from

$$z_{i,j} \sim \text{Multinomial}(\theta_i)$$

- (b) Draw the word based on the chosen topic $z_{i,j}$:

$$w_{i,j} \sim \text{Multinomial}(\phi_{(z_{i,j})})$$

Due to the nature of the topics, they can visually be represented as word clouds, where the size of each word is proportional to its probability of belonging to that topic. However, the topics are not labeled by the procedure, but instead require manual labeling based on the most probable

words per topic. Example of estimated topics and their labeling together with word cloud plots are presented in Appendix B.

We calibrate the procedure to identify 40 topics in the data by maximizing the topics' coherence score C_v (Röder et al., 2015). As topic coherence is generally a measure of the relative distance between words within a topic and thus its interpretability, maximizing this metric is a method for rating the quality of the topics and thereby obtain a set of easily interpreted topics. We then use the estimated conditional probabilities of each word to appear in each topic to decompose each event into the topics it consists of. For example, an article about a meeting between a company's CFO and the minister of commerce of a European country might draw 40% of its words from the topic "international trade", 30% from the topic "corporate governance", and the residual share from other less related topics. These shares represent article-specific topic weights.

We compute firm-specific average topic share $z_{f,\tau,2005}$ by averaging these weights by topic within articles mentioning a specific firm before 2005. Some topics are more common than others. For example, for about 30% of our firms topic 34, relating to trade, is the most common topic among articles mentioning them, and for 90% of our firms it is among the first five topics. Although some topics are very popular overall, there is substantial heterogeneity on which topic firms load the most.

Our hypothesis is that the informational shocks captured in $\tilde{S}_{w(i),f}$ do not have any differential impact on stock market returns on the day of publication of news i , compared to the period preceding the publication, except by affecting the sentiment S_i of a news item published in that week. Therefore we explicitly model our outcome as a difference from its baseline, demeaning an outcome within each event spell. That is, we estimate the effect of the sentiment S_i of news published after 2005-01-01 on $\Delta y_{i,f,n} = y_{i,f,n} - \frac{1}{8} \sum_{k=-10}^{-3} y_{i,f,k}$.

We estimate the model

$$\Delta y_{i,f,n} = \beta S_i + \Lambda_{f,i,n} + \Phi_f + \Phi_{w(i)} + \varepsilon_{i,f,n} \quad (3)$$

$$S_i = \beta^{FS} \tilde{S}_{w(i),f} + \Lambda_{f,i} + \Phi_f + \Phi_{w(i)} + \nu_{i,f} \quad (4)$$

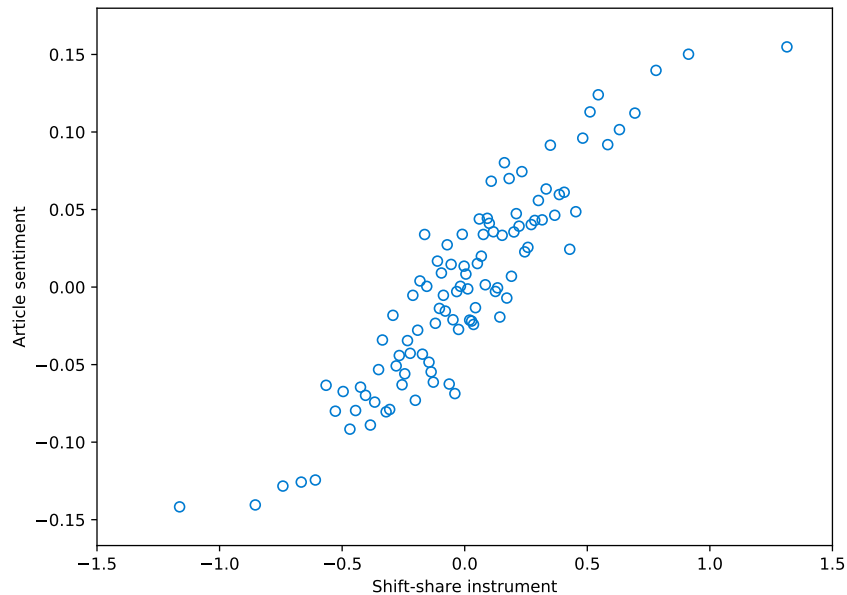
separately for each $n \in -2, -1, 0, 1, 2$ via 2SLS, where $\Lambda_{f,i,n}$ is a vector of control consisting in fixed effects for the day of the week of the observed trading day and firm size. Crucially, as our instrument and source of exogenous variation varies at the week (by news publication time) and firm time, we cluster standard errors at both the week (720 clusters) and firm (680 clusters) level and include fixed effects Φ separately. Therefore, our identification comes exclusively from the interaction of these two terms.

Equation (4) represents the first stage of our instrumental variable model. To interpret our instrumental variable results, we need to additionally assume monotonicity and relevance of the relationship between S_i and $\tilde{S}_{w(i),f}$. Figure 4 plots non-parametrically the conditional relationship between S_i and $\tilde{S}_{w(i),f}$, adjusting for all fixed effects and controls appearing in equations (3) and (4).¹² The figure shows that higher values of $\tilde{S}_{w(i),f}$ predict higher values of S_i . The relationship is robust and significant (with an F-statistics of 134 when tested after the estimation of model (4)).

However, the residual variation in S_i , filtered from idiosyncratic information about a specific firm, does not have any meaningful relationship with abnormal returns at the time of the news

¹²For each figure, we residualize both outcome and predictor. We then group these residuals in fifty equally sized bins ordered according to the predictor residuals, and plot the average within each bin. This approach approximates a Frisch-Waugh-Lovell decomposition.

Figure 4
First stage



NOTE: The figure plots the residuals of a regression of our topic-based shift-share instrument on the fixed effect and controls specified in equations (3) and (4) against the correspondingly computed residuals of our main sentiment score. The residuals are grouped in fifty equally sized bins according to the residuals plotted in the horizontal axis.

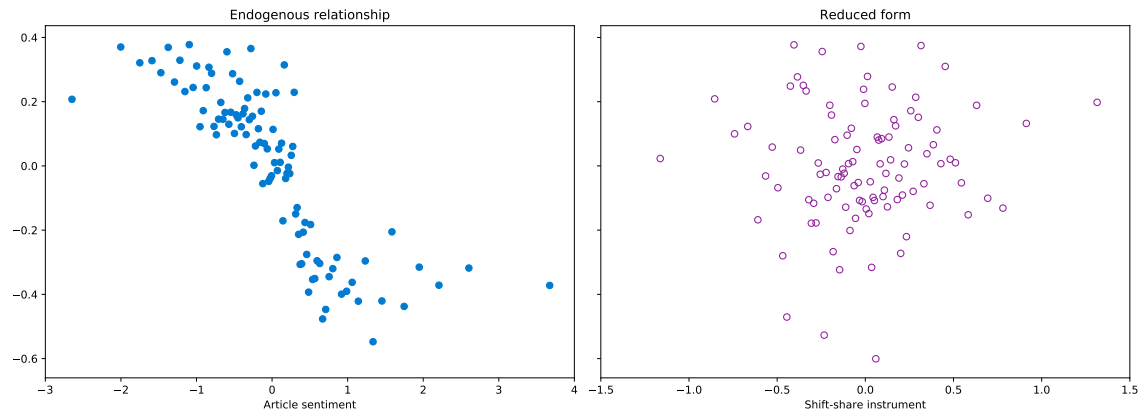
publication. We show the conditional relationship between $\Delta y_{i,f,0}$ and S_i , and between $\Delta y_{i,f,0}$ and $\tilde{S}_{w(i),f}$ in the left and right panels of Figure 5 respectively.

The left panel of Figure 5 shows non-parametrically the results of Section 2. As S_i , which indicates the negativity of an article, increases, abnormal returns decrease with respect to their average before the publication of the news. However, while $\tilde{S}_{w(i),f}$ is strongly related to S_i , there is no discernible association between abnormal returns and our shift-share instrument. The figure suggests that there is no detectable effect of news sentiment on the stock market over and beyond that determined by its idiosyncratic informational content.

Table 4 tests these conjectures formally. We report estimated coefficients for a model estimated for $n \in -1, 0$: the day before and that of news publication. The first and fourth columns report the association between $\Delta y_{i,f,0}$ and S_i , including the variation drive by idiosyncratic news content. These results mirror those obtained in Section 2, although smaller as our sample here only includes articles published in and after 2005.

The second and fifth columns of the table report the reduced form coefficients, obtained by regressing our shift-share instrument directly on the outcome. As the right panel of figure 5 suggests, there is no detectable relationship between our instrument and abnormal returns. As a consequence, the resulting coefficients from the instrumental variable estimation, reported in columns three and six, are extremely noisy, and not statistically distinguishable from zero. We do not find any causal relationship between news sentiment and abnormal returns in the stock market.

Figure 5
The effect of news sentiment on differential abnormal returns at the time of news publication



NOTE: The figure plots the residuals of a regression of our topic-based shift-share instrument and main sentiment score on the fixed effect and controls specified in equations (3) and (4) against the correspondingly computed residuals of the change in abnormal returns at the time of article publication. The residuals are grouped in fifty equally sized bins according to the residuals plotted in the horizontal axis.

The same argument applies to each of our outcomes. Once we disentangle exogenous variation in news sentiment, independent of news content, we find no effect of news sentiment on the stock market. This result implies that if we were to report the same content using optimistic or pessimistic language, the stock market would not be affected. Thus, pessimism cycles do not affect the stock market through the media.

Table 4
The causal effect of news sentiment

| | $n = -1$ | | | $n = 0$ | | |
|-------------------------|-------------------|-------------------|-------------------|-------------------|------------------|------------------|
| | OLS | RF | IV | OLS | RF | IV |
| Abnormal Returns (CAPM) | -0.032 (0.011) | -0.011 (0.032) | -0.070 (0.203) | -0.215 (0.018) | 0.024 (0.051) | 0.154 (0.326) |
| Abnormal Returns (SD) | -0.031 (0.008) | -0.019 (0.025) | -0.119 (0.156) | -0.211 (0.017) | 0.011 (0.044) | 0.073 (0.281) |
| - Own Returns | -0.032 (0.011) | -0.011 (0.032) | -0.071 (0.203) | -0.215 (0.018) | 0.024 (0.051) | 0.153 (0.326) |
| - Sector Returns | -0.002 (0.004) | 0.008 (0.018) | 0.048 (0.113) | -0.003 (0.004) | 0.013 (0.020) | 0.081 (0.125) |
| Book to market | 0.138 (0.067) | 0.326 (0.307) | 2.033 (1.935) | 0.363 (0.111) | 0.492 (0.421) | 3.068 (2.656) |
| # events | 207151 | 207151 | 207151 | 207157 | 207157 | 207157 |
| First stage F-stat | | | 123.872 | | | 123.745 |

NOTE: The table presents the estimated coefficients of the effect of news sentiment on the change in our outcomes one day before and on the day of news publication after 2005. The first and fourth columns reproduce the the results of Section 2, estimating the spurious effect of news sentiment on stock market returns. The second and fifth columns estimate the reduced form effect of our topic-based shift-share instrument on the outcomes. The third and sixth columns estimate the effect of news sentiment, instrumented with our shift-share instrument, on the outcomes of interest. Standard errors double-clustered at the firm and the week of news publication levels are reported in parentheses.

4 Conclusion

This paper isolates variations in firm-specific media coverage sentiment independent of idiosyncratic informational content and the company's stock market performance. We show that while news sentiment is strongly associated with stock market performance, financial markets do not react to exogenous variation in this sentiment. While financial markets care deeply about information, they can see through framing and irrelevant variation in the sentiment of media coverage.

We nonetheless confirm the usefulness of automatically processing news articles in order to capture information transmitted to the market, and replicate earlier results obtained on written newspaper columns in a web-dominated time period. We show that the sentiment of Reuters news about specific companies is strongly associated with the abnormal returns of that company in the short run in the period between 2000 and 2018. However, the relationship is growing weaker over time, and a trading strategy based solely on news sentiment would generate much smaller returns nowadays than it would have done fifteen years ago. This association is approximately 50% stronger during weeks of earning announcements by the company.

References

- AIZER, A. (2010): “The Gender Wage Gap and Domestic Violence,” *American Economic Review*, 100, 1847–1859.
- AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 103, 2121–2168.
- BAKER, M. AND J. WURGLER (2006): “Investor sentiment and the cross-section of stock returns,” *The Journal of Finance*, 61, 1645–1680.
- (2007): “Investor Sentiment in the Stock Market,” *Journal of Economic Perspectives*, 21, 129–152.
- BARTIK, T. J. (1991): *Who benefits from state and local economic development policies?*, W.E. Upjohn Institute.
- BLEI, D. M., A. Y. NG, AND M. I. JORDAN (2003): “Latent Dirichlet Allocation,” *J. Mach. Learn. Res.*, 3, 993–1022.
- BORUSYAK, K., P. HULL, AND X. JARAVEL (2018): “Quasi-Experimental Shift-Share Research Designs,” Working Paper 24997, National Bureau of Economic Research, DOI: 10.3386/w24997.
- BOUDOUKH, J., R. FELDMAN, S. KOGAN, AND M. RICHARDSON (2013): “Which News Moves Stock Prices? A Textual Analysis,” Working Paper 18725, National Bureau of Economic Research, DOI: 10.3386/w18725.
- DELLAVIGNA, S. AND E. KAPLAN (2007): “The Fox News Effect: Media Bias and Voting,” *The Quarterly Journal of Economics*, 122, 1187–1234.
- ENGELBERG, J. E. AND C. A. PARSONS (2011): “The Causal Impact of Media in Financial Markets,” *The Journal of Finance*, 66, 67–97.
- FOGED, M. AND G. PERI (2016): “Immigrants’ Effect on Native Workers: New Analysis on Longitudinal Data,” *American Economic Journal: Applied Economics*, 8, 1–34.
- GERBER, A. S., D. KARLAN, AND D. BERGAN (2009): “Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions,” *American Economic Journal: Applied Economics*, 1, 35–52.
- GOLDSMITH-PINKHAM, P., I. SORKIN, AND H. SWIFT (2018): “Bartik Instruments: What, When, Why, and How,” Working Paper 24408, National Bureau of Economic Research, DOI: 10.3386/w24408.
- HUTTO, C. J. AND E. GILBERT (2014): “Vader: A parsimonious rule-based model for sentiment analysis of social media text,” in *Eighth international AAAI conference on weblogs and social media*.
- JAEGER, D. A., J. RUIST, AND J. STUHLER (2018): “Shift-Share Instruments and the Impact of Immigration,” Working Paper 24285, National Bureau of Economic Research.

- LOUGHRAN, T. AND B. MCDONALD (2011): “When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks,” *The Journal of Finance*, 66, 35–65.
- MCLEAN, R. D. AND J. PONTIFF (2016): “Does academic research destroy stock return predictability?” *The Journal of Finance*, 71, 5–32.
- NAKAMURA, E. AND J. STEINSSON (2014): “Fiscal Stimulus in a Monetary Union: Evidence from US Regions,” *American Economic Review*, 104, 753–792.
- NOVY-MARX, R. (2014): “Predicting anomaly performance with politics, the weather, global warming, sunspots, and the stars,” *Journal of Financial Economics*, 112, 137–146.
- RÖDER, M., A. BOTH, AND A. HINNEBURG (2015): “Exploring the Space of Topic Coherence Measures,” in *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, New York, NY, USA: ACM, WSDM ’15, 399–408.
- TETLOCK, P. C. (2007): “Giving Content to Investor Sentiment: The Role of Media in the Stock Market,” *The Journal of Finance*, 62, 1139–1168.
- (2010): “Does Public Financial News Resolve Asymmetric Information?” *The Review of Financial Studies*, 23, 3520–3557.
- (2011): “All the News That’s Fit to Reprint: Do Investors React to Stale Information?” *The Review of Financial Studies*, 24, 1481–1512.
- TETLOCK, P. C., M. SAAR-TSECHANSKY, AND S. MACSKASSY (2008): “More than Words: Quantifying Language to Measure Firms’ Fundamentals,” *The Journal of Finance*, 63, 1437–1467.
- ZHAO, W., J. J. CHEN, R. PERKINS, Z. LIU, W. GE, Y. DING, AND W. ZOU (2015): “A heuristic approach to determine an appropriate number of topics in topic modeling,” *BMC Bioinformatics*, 16, S8.

A Additional results

Table 5
Dynamic effects, DiD specification

| | P- value | Change in outcome relative to $n = -3$ | | | | | | |
|--|-------------|--|-------------------|-------------------|-------------------|------------------|------------------|------------------|
| | < -3 | -2 | -1 | 0 | 1 | 2 | 5 | 9 |
| <i>Panel A: Loughran/MacDonald composite score</i> | | | | | | | | |
| Abnormal Returns (CAPM) | 0.945 | -0.031 (0.019) | -0.126 (0.023) | -0.501 (0.033) | 0.016 (0.021) | 0.051 (0.020) | 0.050 (0.020) | 0.089 (0.019) |
| Abnormal Returns (SD) | 0.800 | -0.042 (0.015) | -0.136 (0.020) | -0.511 (0.029) | -0.023 (0.017) | 0.022 (0.016) | 0.039 (0.016) | 0.052 (0.016) |
| - Own Returns | 0.945 | -0.031 (0.019) | -0.126 (0.023) | -0.501 (0.033) | 0.016 (0.021) | 0.050 (0.020) | 0.049 (0.020) | 0.089 (0.019) |
| - Sector Returns | 0.822 | 0.011 (0.011) | 0.011 (0.012) | 0.010 (0.011) | 0.039 (0.011) | 0.029 (0.011) | 0.010 (0.011) | 0.037 (0.011) |
| Book to market | 0.315 | 0.122 (0.046) | 0.374 (0.108) | 0.730 (0.146) | 0.711 (0.342) | 0.803 (0.408) | 0.921 (0.413) | 0.847 (0.410) |
| <i>Panel C: First principal component</i> | | | | | | | | |
| Abnormal Returns (CAPM) | 0.431 | -0.014 (0.019) | -0.065 (0.024) | -0.510 (0.031) | -0.010 (0.020) | 0.014 (0.019) | 0.053 (0.019) | 0.078 (0.019) |
| Abnormal Returns (SD) | 0.591 | -0.016 (0.015) | -0.075 (0.021) | -0.509 (0.029) | -0.029 (0.017) | 0.011 (0.016) | 0.032 (0.015) | 0.059 (0.015) |
| - Own Returns | 0.429 | -0.014 (0.019) | -0.065 (0.024) | -0.510 (0.031) | -0.010 (0.020) | 0.014 (0.019) | 0.053 (0.019) | 0.079 (0.019) |
| - Sector Returns | 0.917 | 0.002 (0.011) | 0.010 (0.011) | -0.001 (0.011) | 0.018 (0.011) | 0.003 (0.011) | 0.020 (0.011) | 0.020 (0.011) |
| Book to market | 0.216 | 0.122 (0.056) | 0.304 (0.109) | 0.640 (0.189) | 0.572 (0.308) | 0.663 (0.369) | 0.757 (0.393) | 0.587 (0.412) |

NOTE: The table reports the estimated γ_n coefficients around the time of news publications for a series of stock market outcomes. Standard errors clustered by firm-by-week of publication level are reported in parentheses. The first column reports the p-value of an F-test testing joint equality to zero of all γ_n coefficients for $n < -3$. We can't reject the hypothesis of parallel pre-event trends for any of the outcomes. The first panel uses as a sentiment measure the Loughran/McDonald score, the second, the principal component of the Loughran/McDonald and VADER scores. Both scores are discretized, such that the grouping of the top and bottom 33% of the articles in terms of their negativity scores takes a value of one and zero respectively.

Table 6
Heterogeneity in the effect of sentiment on CAPM abnormal returns, linear specification

| | Number | | P-value | Change in outcome relative to $n = -3$ | | | | |
|--------------------------------------|--------|----------|---------|--|-------------------|-------------------|-------------------|-------------------|
| | events | clusters | < -3 | -2 | -1 | 0 | 1 | 2 |
| <i>Panel A: Over sectors</i> | | | | | | | | |
| Communication Services | 22697 | 7426 | 0.277 | 0.001 (0.025) | -0.056 (0.026) | -0.143 (0.034) | -0.039 (0.025) | -0.037 (0.025) |
| Consumer Discretionary | 35399 | 16819 | 0.471 | -0.000 (0.024) | -0.055 (0.029) | -0.419 (0.040) | -0.002 (0.026) | 0.016 (0.025) |
| Consumer Staples | 13727 | 7238 | 0.153 | 0.036 (0.030) | -0.014 (0.031) | -0.314 (0.059) | 0.018 (0.031) | 0.052 (0.028) |
| Energy | 23497 | 10776 | 0.500 | -0.022 (0.032) | -0.034 (0.040) | -0.126 (0.038) | -0.016 (0.037) | 0.018 (0.035) |
| Financials | 42247 | 18327 | 0.783 | 0.014 (0.025) | -0.068 (0.034) | -0.148 (0.037) | -0.008 (0.026) | 0.012 (0.027) |
| Health Care | 34653 | 15238 | 0.811 | 0.015 (0.018) | -0.064 (0.021) | -0.268 (0.031) | -0.015 (0.018) | 0.003 (0.017) |
| Industrials | 34320 | 14503 | 0.807 | -0.011 (0.024) | -0.021 (0.024) | -0.242 (0.036) | 0.020 (0.026) | 0.032 (0.025) |
| Information Technology | 41846 | 16095 | 0.797 | -0.029 (0.022) | -0.064 (0.024) | -0.432 (0.040) | -0.008 (0.022) | 0.016 (0.022) |
| Materials | 15562 | 7678 | 0.117 | -0.080 (0.035) | -0.120 (0.043) | -0.257 (0.047) | -0.053 (0.035) | -0.096 (0.035) |
| Real Estate | 4107 | 2760 | 0.435 | 0.027 (0.074) | 0.057 (0.059) | 0.021 (0.073) | 0.203 (0.070) | 0.100 (0.061) |
| Utilities | 17012 | 8010 | 0.346 | 0.012 (0.033) | -0.075 (0.041) | -0.070 (0.040) | 0.033 (0.045) | 0.073 (0.041) |
| <i>Panel B: Earning announcement</i> | | | | | | | | |
| No | 153311 | 73080 | 0.943 | -0.005 (0.010) | -0.050 (0.013) | -0.208 (0.016) | -0.001 (0.011) | 0.008 (0.011) |
| Yes | 131756 | 53927 | 0.726 | -0.006 (0.014) | -0.059 (0.016) | -0.306 (0.021) | -0.016 (0.015) | 0.013 (0.015) |

B Latent Dirichlet Allocation for Topic Identification

For preprocessing the news articles we follow standard procedures: We first remove punctuation marks, newlines and tabs and convert to lowercase. Then we remove stopwords (such as *the*, *is*, *are*, and *this*) and lemmatize all words, where the purpose of the latter is to reduce words to their respective word stems in order to limit the textual variability across documents. Finally, we trim the corpus such that tokens that occur less than 15 times and in more than 50% of the documents are removed in order to filter tokens that are either very rare or typical. This procedure returns a final dictionary with approximately 100 000 tokens. A few estimated topics and their labeling are shown in Table 7.

Table 7
Selected estimated topics and their labeling

| Topic ID | Label | Keywords |
|----------|-------------|--|
| 0 | Energy | production, barrel, company, oil , gas , energy, pipeline, crude, drill |
| 16 | Germany | germany , european, euros , german , euro, europe |
| 19 | Automobiles | sales , motor , ford , car, model, market, vehicles, industry, auto, cars |

NOTE: The table reports the estimated topics and their labeling from an LDA-procedure applied to the Reuters news corpus. Topics are labeled based on the most important words marked in bold.

DANMARKS NATIONALBANK
HAVNEGÅDE 5
DK-1093 COPENHAGEN K
WWW.NATIONALBANKEN.DK



As a general rule, Working Papers are not translated, but are available in the original language used by the contributor.

Danmarks Nationalbank's Working Papers are published in PDF format at www.nationalbanken.dk. A free electronic subscription is also available at this Website. The subscriber receives an e-mail notification whenever a new Working Paper is published.

Text may be copied from this publication provided that Danmarks Nationalbank is specifically stated as the source. Changes to or misrepresentation of the content are not permitted.

Please direct any enquiries to Danmarks Nationalbank, Communications, kommunikation@nationalbanken.dk