

Media abnormal tone, earnings announcements, and the stock market

D. Ardia, K. Bluteau, K. Boudt

G-2021-50

September 2021

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

Citation suggérée : D. Ardia, K. Bluteau, K. Boudt (Septembre 2021). Media abnormal tone, earnings announcements, and the stock market, Rapport technique, Les Cahiers du GERAD G- 2021-50, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2021-50>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Suggested citation: D. Ardia, K. Bluteau, K. Boudt (September 2021). Media abnormal tone, earnings announcements, and the stock market, Technical report, Les Cahiers du GERAD G-2021-50, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (<https://www.gerad.ca/en/papers/G-2021-50>) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2021
– Bibliothèque et Archives Canada, 2021

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2021
– Library and Archives Canada, 2021

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Télec. : 514 340-5665
info@gerad.ca
www.gerad.ca

Media abnormal tone, earnings announcements, and the stock market

David Ardia ^a

Keven Bluteau ^b

Kris Boudt ^{c, d, e}

^a GERAD & Department of Decision Sciences,
HEC Montréal, Montréal (Qc), Canada, H3T 2A7

^b Department of Finance, Université de Sherbrooke,
Sherbrooke (Qc), Canada, J1K 2R1

^c Department of Economics, Ghent University,
9000 Ghent, Belgium

^d Solvay Business School, Vrije Universiteit Brussel,
1050 Brussels, Belgium

^e School of Business and Economics, Vrije Universiteit
Amsterdam, 1081 HV Amsterdam, The Netherlands

david.ardia@hec.ca

keven.bluteau@usherbrooke.ca

kris.boudt@ugent.be

September 2021
Les Cahiers du GERAD
G–2021–50

Copyright © 2021 GERAD, Ardia, Bluteau, Boudt

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract : We propose a tone-based event study to reveal the aggregate abnormal tone dynamics in media articles around earnings announcements. We test whether they convey incremental information that is useful for price discovery for non-financial S&P 500 firms. The positive relationship found between the abnormal tone and abnormal returns suggests that media provide incremental information relative to the information contained in earnings press releases and earnings calls.

Keywords: Abnormal returns, abnormal tone, earnings announcements, event study, news media, sentometrics

Acknowledgements: We thank the Editor of Journal of Financial Markets (Tarun Chordia); an anonymous referee; Francesco Audrino; Marie-Claude Beaulieu; Samuel Borms; Tolga Cenesizoglu; Serge Darolles; Kilian Dinkelaker; Michel Dubois; George Hübner; Alexandre Jeanneret; Tim Kroencke; Marie Lambert; Gaëlle Lefol; Louis Mangeney; Hannes Mohrschladt; Carlos Ordás Criado; Christophe Perignon; Jérôme Taillard; James Thewissen; Philip Vervliet; seminar participants at Florence University, HEC Montréal, Skema Business School, and the University of Delaware; as well as participants at AFFI (Nantes, 2021, and Québec, 2019), CFE (Pisa, 2018, and London, 2019), and R/Finance (Chicago, 2018) conferences for helpful comments. We acknowledge the Flemish Science Foundation (<https://fwo.be>), Innoviris (<https://innoviris.brussels>), IVADO (<https://ivado.ca>), and the Swiss National Science Foundation (<https://www.snf.ch>, grants #179281 and #191730) for their financial support.

1 Introduction

An earnings press release provides investors with new information about a company. This happens directly, through the publication of earnings numbers, as well as indirectly, by triggering qualitative discussions of a firm's prospects, such as through earnings calls and media reports. Current evidence suggests that the media is a double-edged sword for financial markets. On the one hand, the media can reduce information asymmetry by quickly and broadly distributing relevant information, leading to a more efficient market (see, e.g., Fang and Peress, 2009; Tetlock, 2010). On the other hand, the media can produce short-term price distortions, as news reports can exacerbate the effects of investor biases (see, e.g., Tetlock, 2007; Dougal et al., 2012; Garcia, 2013; Shiller, 2017).

This paper uses a novel tone-based event study to evaluate how the media report around earnings announcements and how incremental information provided by the media can help explain stock market reactions around earnings announcements. To put these research questions into context, Figure 1 displays a flowchart illustrating in a stylized way the information flow between firms, the media, and investors. According to the Regulation Fair Disclosures (Reg FD), publicly traded firms are obliged to disclose all material non-public information. In case of earnings announcements, this is done via earnings calls and earnings press releases, as illustrated in the “primary information” layer of Figure 1. Through those channels, firms provide media and investors with quantitative and qualitative information about their financial results and business prospect.

When the information is of interest to their readers, news media outlets then report about the earnings event, often contextualizing the information and adding analysts' comments and forecasts. As formalized by Nimark and Pitschner (2019), media thus act as information intermediaries between investors and the true state of the world of the firm (see also Larsen, Thorsrud, and Zhulanova, 2021). The second layer in Figure 1 relates to this indirect transmission, where media disseminate and transform the first layer of information to investors. Under this framework, we refer to dissemination when the media transmit the most important information in the first layer, whereas transformation corresponds to any modification, completion, or contextualization of the information available in that layer.

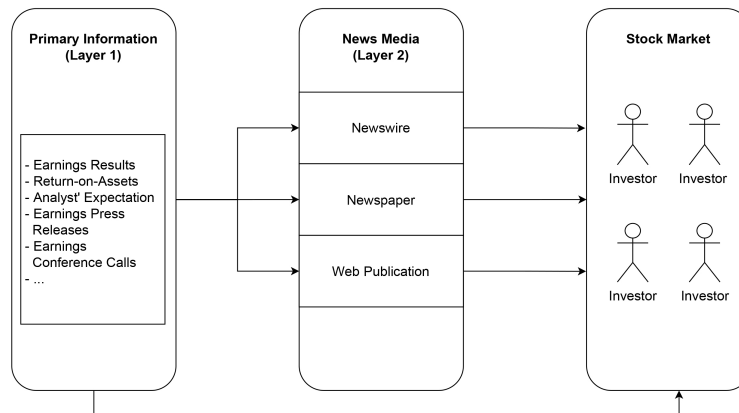


Figure 1: Information flow. This chart illustrates in a stylized way the information flow between firms, media, and investors during an earnings announcement event.

As a concrete example of the dissemination and transformation of the media, consider the earnings announcement of Home Depot on August 19, 2008. The firm reports earnings per share of 71 cents that beat the average analyst forecast of 61 cents. One of the news articles in our corpus reporting about that event is published by CNN with the title “Home Depot profit tumbles 24%.” In that article, the journalist disseminates the information from the earnings press release and earnings call about the higher-than-expected earnings and its relation to the stimulus checks sent by the government in May. Moreover, the journalist transforms the first layer of information by referring to another article

published by CNN stating that “(...) sales dropped off in July, suggesting the program had run its course.” Hence, the increase explained by the stimulus checks is not deemed sustainable.¹

The dissemination and transformation of the media are reflected in the tone. Based on the above framework, we posit two hypotheses. First, the media tone provides a complementary signal about the firm compared to what can be predicted using only the information derived from the earnings report, the earnings press release, and the earnings call. The complementary role of news media stems from their transformation of earnings-related information into a more easily understandable, contextualized, and condensed format, as well as the additional information it can provide given important contemporary events. Second, controlling for the information contained in earnings press releases and earnings calls, we expect that the tone of media articles is useful for stock price discovery following an earnings announcement. This hypothesis implies that the media contributes to disseminating and transforming information to investors during earnings announcements.

To test this empirically, we introduce a novel tone-based event study. We generalize the data-driven word power methodology of Jegadeesh and Wu (2013) to the setting of multiple news articles per day to compute the daily media tone for firms.² Then, we proceed in two steps. First, for a given firm and time, we define an “abnormal tone” as an unexpected tone conditional on a market capitalization-weighted average of the media tone about firms in the universe. The abnormal tone is designed to capture the idiosyncratic part of the news as opposed to news generated by general market conditions. We aggregate the abnormal tone into a “cumulative abnormal tone” (CAT) measure around the event. Then, we filter the CAT to remove the information that can be derived from earnings press releases and earnings calls, as well as other variables such as book-to-market ratio and firm size. We refer to this incremental information measure as the “residual cumulative abnormal tone” (RCAT).³ We apply our approach to media articles on non-financial S&P 500 firms’ quarterly earnings announcements from 2000 to 2016. In doing so, we analyze more than 2.3 million news articles published by newswires, newspapers, and web publications.

We first analyze the drivers of CAT around earnings announcements. As expected, we find that CAT is largely driven by the information disclosed by the firm, such as the sign of the earnings surprise, the tone of the press release, and the tone of the earnings call. CAT is also driven by its past values, suggesting that the new information embedded in CAT also incorporates past information. However, a large part of the CAT’s variance remains unexplained, which supports the use of RCAT to analyze incremental information from media tone.

We then investigate if RCAT conveys incremental information about a firm’s stock price reaction around the announcement. We find a significant positive contemporaneous relationship: an increase to RCAT of one standard deviation is associated with an expected excess abnormal return of about 0.84%. Furthermore, we find that RCAT has *predictive* power for future price movements. Specifically, we find a partial reversal in stock price movements associated with RCAT in the month following the announcement; a one standard deviation increase in RCAT is associated with an excess abnormal return of -0.36%. One possible channel is that stock markets overreact to incremental information conveyed by the media’s abnormal tone. This interpretation is consistent with our finding that the absolute value of RCAT is positively related to a firm’s abnormal trading volume around the earnings announcement.

¹The two news article published by CNN are available at https://money.cnn.com/2008/08/19/news/companies/home_depot_earnings/index.htm?postversion=2008081906 and https://money.cnn.com/2008/08/13/news/economy/retail_sales/index.htm?postversion=2008081311, the earnings press release at <https://ir.homedepot.com/news-releases/2008/08-19-2008>, and the earnings call transcript at <https://seekingalpha.com/article/91649-home-depot-f2q08-qtr-end-8-3-08-earnings-call-transcript?part=single>.

²Jegadeesh and Wu (2013) quantify the tone of words in 10-Ks by means of a regression of the firm’s stock return on the presence of sentiment words in the corresponding 10-K. They thus have one document per observation period while we have multiple.

³Huang, Teoh, and Zhang (2013) pioneer the approach of analyzing incremental tone by means of the residuals of a regression. They do this for earnings press releases and interpret the residual tone as the tone management by the firm. It is inspired by the calculation of discretionary accruals in accounting and finance.

This article contributes to two different strands of research. First, it advances the literature on the *quantification* of information available in textual data to better understand the role of the news media in financial markets. Early pioneers in this literature include Tetlock, Saar-Tsechansky, and Macskassy (2008) and Engelberg (2008), who all used the General Inquirer (Stone and Hunt, 1963) system to compute the sentiment of media publications such as Wall Street Journal columns and newswire articles. Since Loughran and McDonald (2011), it has become best practice to use a domain-specific lexicon in accounting and finance. While several lexicons exist to link the stock price reaction to the textual information in corporate disclosures, there is no such lexicon for media analysis. As also advocated by Ke, Kelly, and Xiu (2019), the use of a data-driven calibration of a lexicon fills that gap. In this study, we propose an extension of the word power methodology proposed by Jegadeesh and Wu (2013) to obtain a domain-specific lexicon for analyzing the daily tone of firms in media articles with the goal of explaining the stock market. We then develop a novel and comprehensive event-study framework designed to isolate the information from the daily tone that is incremental to other sources of information. We further use a structural topic model (STM) approach (Roberts, Stewart, and Airoldi, 2016) to analyze the difference in content between the three types of sources: media news articles, earnings press releases, and earnings calls. Compared to the traditional topic model, the STM has the advantage that the effects of metadata can be modelled. We exploit this feature to provide insights on how media differentiate themselves from earnings press releases and earnings calls by focusing on certain subjects. As such, we go beyond the analysis of tone in studying the mechanism through which the media provide additional information to the market. Such a topic analysis provides insights on how media focus on certain subjects of discussion to contextualize, emphasize, or mitigate specific results compared to earnings press releases and earnings calls. While the focus of our study is on earnings announcements, our analysis can easily be transposed to analyze the role of media in transforming the information around other events, such as central bank announcements, natural disasters, or other corporate events.

Second, the results contribute to the vast literature on price discovery around corporate disclosure events, and in particular, earnings announcements. Several studies have focused on the information contained in the tone of newswire articles (see Engelberg, 2008), the tone of earnings press releases (see Henry, 2008; Demers and Vega, 2008; Davis, Piger, and Sedor, 2012; Huang, Teoh, and Zhang, 2013; Arslan-Ayaydin, Boudt, and Thewissen, 2016), the tone of earnings calls (see Price et al., 2012; Borochin et al., 2018), the narrative structure of voluntary disclosure narratives (see Allee and DeAngelis, 2015), as well as their level of readability (see Jones and Shoemaker, 1994; Li, 2008; Merkley, 2014; Asay, Libby, and Rennekamp, 2018). Whether markets overreact or underreact to the abnormal tone in media is still an open question. Engelberg (2008) analyzes only the Dow Jones Newswires publications and finds an underreaction that is only economically significant for the 40 to 80 days following the announcement.⁴ The literature on the investors' attention bias would rather predict an overreaction to the tone of media articles that capture the attention of retail investors (see, e.g., Da, Engelberg, and Gao, 2011). We test the two hypotheses using a comprehensive corpus of news from several newswire, newspaper, and web publication sources. Our findings confirm that there is an overreaction to the abnormal tone of media news articles leading to a price reversal in the 20 days following the event.

2 Tone-based event study

Algaba et al. (2020) define tone as “the disposition of an entity toward an entity expressed via a certain medium.”⁵ This disposition can originate from facts (i.e., “the firm’s EPS beats analyst consensus”), referred to as *factual sentiment* by Heerschoep et al. (2011), or opinions (i.e., “the firm announced a

⁴See Figure 2 in Engelberg (2008).

⁵Other popular terms are “sentiment,” “polarity,” or “opinion.” Under the tone implementation of our paper, the tone can be interpreted as the “expected return implied by textual documents.” Throughout the paper, we use “tone” for two reasons. First, it is the naming convention used in Jegadeesh and Wu (2013), which is a reference in the field. The authors use the word “tone” more than a hundred times in the context of 10-Ks, which are also factual documents.

disappointing quarter, but discussed an innovative and impressive new product line”). Kim (2019), for instance, differentiate between *factual tone* and *promotional tone*. Given the difficulty of identifying facts from opinions in an automated way, we do not differentiate between them in this paper.

Below, we present our approach to extract the (abnormal) tone about firms from a sample of media articles published close to earnings announcements (regardless of whether the articles discuss earnings results). Our implementation follows Jegadeesh and Wu (2013) and leads to tone estimates that can be interpreted as “the expected stock return implied by the textual documents.” While this follows strictly the naming convention used by Jegadeesh and Wu (2013), we wish to emphasize that the proposed event-study approach is agnostic of the tone computation method. It can be implemented with various alternative methods (supervised or not) than the approach of Jegadeesh and Wu (2013).

2.1 Measurement of the tone

We follow Jegadeesh and Wu (2013) in using a data-augmented bag-of-words approach to measure the tone in news articles about firms. The calibration methodology of the lexicon, detailed in the Internet Appendix, Section A, leads to a dictionary of J words with associated polarity scores ζ_j , $j = 1, \dots, J$, where the scores set the polarity of each word (i.e., the degree of positivity or negativity). For a firm $k = 1, \dots, K$, the tone at time t is defined as the weighted average of the polarity scores of the words in the text:

$$TONE_{k,t} = \sum_{j=1}^J \zeta_j f(j, k, t), \quad (1)$$

where $f(j, k, t)$ maps a word j for firm k at time t to a real number.

We define the function f as an average of term frequencies in news articles published at time t . More precisely, let $D_{k,t}$ be the number of articles written about firm k at time t , conditional on the presence of at least one word j in each article $d = 1, \dots, D_{k,t}$, $FQ_{d,j,k,t}$ as the number of times the j th word is encountered in article d , and $N_{d,k,t}$ as the number of words in article d . Then, the average term-frequency for word j , firm k , at time t is:

$$f(j, k, t) = \frac{1}{D_{k,t}} \sum_{d=1}^{D_{k,t}} \frac{FQ_{d,j,k,t}}{N_{d,k,t}}. \quad (2)$$

The word power definition of $f(j, k, t)$ from Jegadeesh and Wu (2013) is a special case of our definition where $D_{k,t} = 1$ at all times for all firms. Due to the flexibility in defining $f(j, k, t)$, we refer to our methodology as the generalized word power (GWP).

2.2 Definition of the cumulative abnormal tone

Given an event-firm i and firm-related articles published at time τ relative to the earnings announcement date t_i , we decompose tone into the sum of a “normal tone” and an “abnormal tone”:

$$TONE_{i,\tau} = NTONE_{i,\tau} + ATONE_{i,\tau}, \quad (3)$$

where the relative time τ is zero for the earnings announcement day. We use this decomposition of the observed tone to extract the new information as revealed in the abnormal tone. This interpretation is motivated by the fact that, in the hypothetical case where media predominantly publish information already known or expected, the abnormal tone should be much closer to zero than if the media just published new information.⁶

Second, our event-study framework can be used with alternative methods (supervised or not) than the word power to extract the tone, which do not necessarily lead to an interpretation as “expected return implied by textual documents.”

⁶Huang, Teoh, and Zhang (2013) and Arslan-Ayaydin, Boudt, and Thewissen (2016) use a measure of abnormal tone for firms’ earnings press releases where the normal tone is modeled using firms’ fundamentals. They define the abnormal tone as a measure intended to capture the discretionary and inflated component of tone. We later use this method to extract the media’s transformation of information about earnings announcements.

Finally, to evaluate the effect of an earnings announcement on the abnormal tone over a time window of τ_1 to τ_2 ($\tau_1 < \tau_2$) relative to the event day, we define the cumulative abnormal tone, CAT, as follows:

$$CAT_i(\tau_1, \tau_2) = \sum_{\tau=\tau_1}^{\tau_2} ATONE_{i,\tau}. \quad (4)$$

2.3 Estimation of the normal tone model

The decomposition of tone into normal and abnormal components in (3) requires an estimate of the normal tone. For this purpose, we use a linear factor model:

$$TONE_{i,\tau} = \alpha_i + \beta_i F_\tau + \epsilon_{i,\tau}, \quad (5)$$

where F_τ is a tone factor common across all firms at relative time τ ; α_i and β_i are the event-firm constant and factor exposure, respectively; and $\epsilon_{i,\tau}$ is an error term. This specification is analogous to the model of expected returns commonly used in abnormal-return event studies (see, e.g., MacKinlay, 1997).⁷ For the tone factor F_t , we use a weighted average of the firms' tones for the K_t firms for which we can compute a tone value at time t :

$$F_t = \sum_{k=1}^{K_t} \omega_{k,t} TONE_{k,t}, \quad (6)$$

where $\omega_{k,t}$ is the firm- k weight at time t and $\sum_{k=1}^{K_t} \omega_{k,t} = 1$. F_τ is then defined as the appropriately aligned value of F_t at relative time τ .

We estimate the constant α_i and the factor exposure β_i in (5) by ordinary least squares using data available in an estimation window $[t_i - L - K, \dots, t_i - K - 1]$, where L is the length of the window and K is the offset relative to the event date t_i . The abnormal tone is then measured over an event window:

$$ATONE_{i,\tau} = TONE_{i,\tau} - (\hat{\alpha}_i + \hat{\beta}_i F_\tau), \quad (7)$$

for $\tau > t_i - K - 1$, where $\hat{\alpha}_i$ and $\hat{\beta}_i$ are the estimated event-entity-specific constant and factor exposure, respectively.

For illustration, Figure 2 shows how our methodology applies to Home Depot's 2008 second-quarter earnings announcement (August 19, 2008). First, we collect all relevant news articles published about Home Depot for days around August 19, 2008. Second, for each day, we measure the daily tone extracted from the news articles using the GWP model calibrated on data observed up to December 31, 2007. Third, we estimate the normal tone model with $L = 30$ and $K = 5$ days. Finally, we compute the abnormal tone and CAT over various post-estimation windows of interest (pre-event, event, and post-event windows).

2.4 Definition of residual CAT

Our definition of abnormal tone captures the idiosyncratic part of the news as opposed to news generated by general market conditions. We can go one step further and isolate or filter out a specific part of the new information. In particular, in the case of scheduled news such as earnings announcements, we can remove the effect of the firm's communication and financial results from the media's reporting to assess how the media transforms that information. To do so, we follow the procedure outlined in Huang, Teoh, and Zhang (2013). Specifically, we define the residual CAT of event-firm i , $RCAT_i(\tau_1, \tau_2)$, by the transformation of $CAT_i(\tau_1, \tau_2)$ such that it is orthogonal to a vector \mathbf{x}_i of information originating from first-layer sources (e.g., the firm and market participants):

$$RCAT_i(\tau_1, \tau_2) = CAT_i(\tau_1, \tau_2) - \hat{\boldsymbol{\theta}}'_{\tau_1, \tau_2} \mathbf{x}_i. \quad (8)$$

⁷This model nests the special case of setting the normal tone to an event-specific constant ($\beta_i = 0$).

The vector of coefficients $\hat{\theta}_{\tau_1, \tau_2}$ is estimated from the regression $CAT_i(\tau_1, \tau_2) = \theta'_{\tau_1, \tau_2} \mathbf{x}_i + \eta_i$, where η_i is an error term. Primary information for earnings announcements include the tone of the earnings press release, the tone of the earnings call, the earnings surprise, and the return on assets.

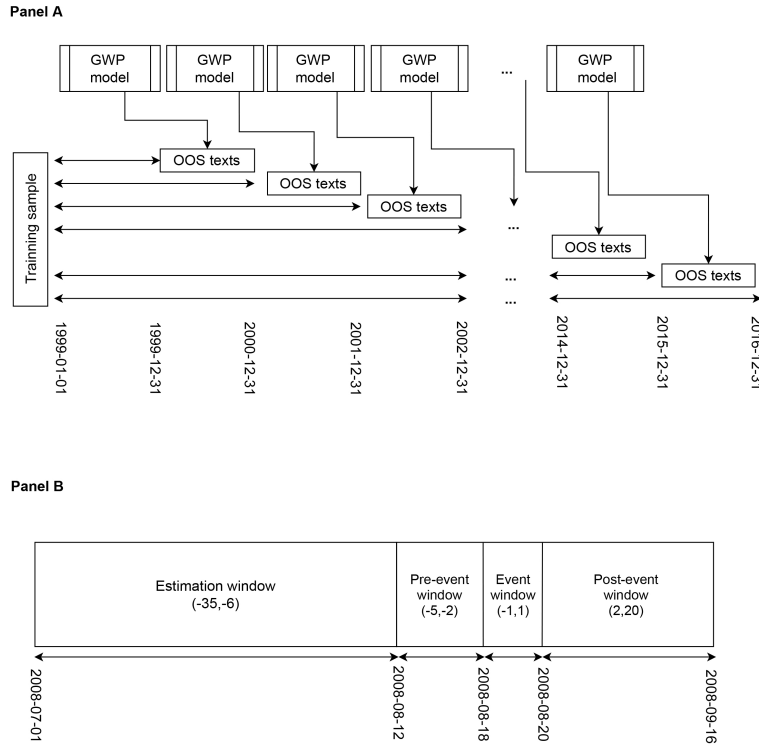


Figure 2: Abnormal tone event study timing information. This figure shows the GWP tone estimation scheme (Panel A) and the event study methodology (Panel B). The GWP model is trained at the end of each year using an expanding window (with the sample data range indicated on the horizontal lines). Then, each GWP-trained model is used to compute the tone of news articles for the following year (OOS text boxes). Panel B is based on the Home Depot earnings announcement event for the second quarter of 2008. We use an estimation window of $L = 30$ days, with an offset of $K = 5$ days, and an event window of 20 days. “TD” stands for trading days.

2.5 Residual and abnormal tone of individual news articles

To investigate a specific news article (or a group of articles), we isolate the contribution of each article to the daily (abnormal) tone. The tone of a single article d discussing event-firm i and published at relative time τ is obtained as:

$$TONE_{d,i,\tau} = \sum_{j=1}^J \hat{\zeta}_j \frac{FQ_{d,j,i,\tau}}{N_{d,i,\tau}}, \quad (9)$$

from which we can compute the abnormal tone of a specific news article as:

$$ATONE_{d,i,\tau} = TONE_{d,i,\tau} - (\hat{\alpha}_i + \hat{\beta}_i F_\tau). \quad (10)$$

Finally, we can measure the article-specific residual abnormal tone contribution, $RATC_{d,i,\tau}(\tau_1, \tau_2)$, as follows:

$$RATC_{d,i,\tau}(\tau_1, \tau_2) = \frac{ATONE_{d,i,\tau}}{D_{i,\tau}} - \frac{\hat{\theta}'_{\tau_1, \tau_2} \mathbf{x}_i}{\sum_{j=\tau_1}^{\tau_2} D_{i,j}}, \quad (11)$$

where $\tau_1 \leq \tau \leq \tau_2$.

3 Data

Our empirical analysis focuses on news articles discussing 598 non-financial firms that were included in the S&P 500 index between 2000–Q1 and 2016–Q4.⁸

3.1 Earnings, accounting, and stock market data

We collect quarterly earnings dates, values, and analyst forecasts from the I/B/E/S database.^{9,10} We gather the quarterly asset value, net income, book value, market capitalization, daily stock prices, and trading volumes from the merged CRSP-Compustat database. We match the two data sources using the I/B/E/S CUSIP and the Compustat NCUSIP.

To identify “good” and “bad” earnings announcements, we follow Livnat and Mendenhall (2006) and compute the analysts’ standardized unexpected earnings (also called the “earnings surprise”):

$$SUE_i = \frac{EPS_i - MFOR_i}{P_i}, \quad (12)$$

where, for event-firm i , EPS_i is the reported earnings per share, $MFOR_i$ is the median analyst forecast of earnings per share, and P_i is the price of the firm at the end of the earnings quarter. We consider only the most recent forecast for each analyst produced no more than 90 days before an earnings announcement.

To measure the share turnover, we follow Campbell and Wasley (1996) and compute the log-percentage of outstanding shares traded:

$$ST_{i,\tau} = \log \left(\frac{V_{i,\tau}}{C_{i,\tau}} \right), \quad (13)$$

where $V_{i,\tau}$ is the trading volume and $C_{i,\tau}$ is the number of common shares outstanding, for event-firm i at relative time τ .

3.2 Textual data

First, we collect firms’ historical company names and tickers from Compustat. For each historical company name, we retrieve all news articles available on LexisNexis with a minimum relevance score of 85.¹¹ This database is used, for instance, in Fang and Peress (2009) and Ahmad et al. (2016). Our news sample spans from January 1, 1999, to December 31, 2016.¹²

We then use the following filters:

- We only keep news articles categorized by LexisNexis English as a “newswire,” “newspaper,” or “web publication.”
- We eliminate articles with fewer than 200 words, as tone measures for short articles are noisy (see Shapиро, Sudhof, and Wilson, 2017).

⁸We only consider firms at the time they were included in the S&P 500 index. Our sample thus tracks S&P 500 constituents over time. We exclude the financial sector, as is commonly done in earnings-announcement event studies. Financial firms are identified using the first two digits of the Global Industry Classification Standard code for each firm.

⁹We obtain forecasts of quarterly figures reported in the quarter prior to a quarterly report’s publication (i.e., $FPI=6$) from the Forecast Period Indicator.

¹⁰Earnings announcements can be made outside of trading hours. As such, there can be differences in the timing of market reaction to earnings announcements. To account for this, we follow Engelberg, McLean, and Pontiff (2018), and redefine each earnings announcement date as the day with the highest trading volume in a three-day window centered on the I/B/E/S earnings announcement date.

¹¹LexisNexis indexes each article with metadata information such as the company or topic of the article. These metadata tags are associated with a relevance score indicating if there is a major reference to the tag. To avoid sampling errors, we manually verify that each historical company name is matched to the LexisNexis equivalent company metadata tag and remove company names without a LexisNexis equivalent company tag.

¹²We use 1999 data to calibrate the word scores for 2000.

- We exclude patent announcements by leveraging LexisNexis' topic classification, as they are irrelevant to our study.
- We remove duplicate (or near-duplicate) articles using locality-sensitive hashing (see Wang et al., 2014).
- We remove machine-generated texts, such as automatic daily stock picks from newswires. These articles are mainly composed of numbers, are highly structured, and lack opinions because they are not written by individuals.
- We remove articles for which there are more than two firms (or tickers) with a major reference (i.e., relevance score of at least 85) to avoid mixing the tone of several firms (see Kelley and Tetlock, 2013).
- We remove the sources "Midnight Trader - Live Briefs" and "News Bites - US Markets," which have a large number of news items (they are in the top 100 sources based on the number of news items), but predominantly provide market summaries and trading recommendations. In most cases, the type of news articles provided by these sources are automatically generated.¹³

Each news article is processed to remove web address (URL), numbers, punctuation, spaces, and other non-textual material. We then transform all words to their root form using the Porter stemming algorithm (Porter, 1980). More details on the various algorithms and software used in the cleaning process are provided in the Internet Appendix, Section B.

The resulting database contains earnings press releases and transcripts of the earnings calls (distributed through newswires), and news articles relevant to test our hypotheses on media transformation. It is crucial to split news incoming from the media and news incoming from the firm to evaluate the relationship between the media and the market truly. We rely on rule-based procedures to identify the earnings press releases and the earnings call transcripts in the news database. For press releases, we extract articles from newswires that have been tagged as either "press releases," "company earnings," "financial results," "interim financial results," "financial performance and reports," or "earnings per share." We also require that the articles contain the word "contact," as contact information is typically included in earnings press releases. We then apply a negative filter over the selected newswire articles using keywords associated with misclassifying newswire articles as earnings press releases. Finally, in case of multiple detections of earnings press releases in the newswire data, we select the most likely one based on keyword similarity with a training set of earnings press releases. For earnings calls transcripts, we extract documents classified as "transcript" in LexisNexis for each firm. Again, we use keyword similarity selection in case of multiple detections.

In Figure 3, we display the number of news articles retrieved by publication type over time. Newswires dominate other sources for the entire sample. We also see substantial growth in the number of articles over time. This is due to an increase in the publication frequency and the number of sources available. For example, the number of articles published about S&P 500 firms originating from web publications is negligible before 2008 but dominates newspapers after 2010.

In Table 1, we report summary statistics on the number of articles and daily coverage. We define the daily coverage as the percentage of trading days with at least one news item discussing the firm. Overall, our sample contains 2,319,582 news articles, with an average of 3,878 articles per firm and an average daily coverage of 34.85% (i.e., for a typical firm, on about one in every three trading days, there will be at least one article published about the firm; on the other two trading days, there will be no articles). The minimum daily coverage for a firm in our sample is 0.27%, and the maximum daily coverage is 100%. We further split the sample across four buckets of 149 firms conditional on the average market capitalization computed for each firm over the 1999 to 2016 period. Firms with the lowest market capitalization have an average of 948 news articles discussing them and an average daily coverage of 20.67%. Firms with the highest market capitalization have an average of 10,449 news articles and an average daily coverage of 54.82%. Hence, there is a positive relationship between the

¹³We verified that the conclusions of this study remain qualitatively similar when we remove the last two filters.




Figure 3: Number of articles over time. This figure shows the number of news articles per year by source type from 1999 to 2016.

number of articles (and the daily coverage) and the market capitalization of the firm (at least in the S&P 500 universe). This relationship holds across different types of publications.¹⁴

In the Internet Appendix, Section C, we list the top sources of news articles, for which we report the source name, the source type, the number of news articles included in our sample, and a brief description of the source and target content.

3.3 Tone measurement for news articles

We now describe how we calibrate our GWP methodology to our data set. To capture the market's reaction to information published by the media, we follow Jegadeesh and Wu (2013) and calibrate the word scores on firms' stock returns. This way, we can also interpret the computed textual tone as the expected return implied by the media articles. The list of words in our dictionary is obtained by merging the Loughran and McDonald (2011) and Harvard IV-4 (Stone and Hunt, 1963) lexicons. We use the Porter stemming algorithm (Porter, 1980) to eliminate noninformative word variations. Our lexicon contains 3,585 root words.

We train the GWP and measure the tone using an expanding window (see Panel A of Figure 2). To ensure that the market reaction is concurrent with news articles, we only use newswires to calibrate the GWP, as newspapers and web publications can be delayed (see Fang and Peress, 2009). In addition, to obtain reliable score estimates, we remove words that are observed in fewer than 200 days across all firms for each expanding window. Finally, to avoid any look-ahead bias, we use the calibrated scores to measure the tone of the news articles of the year following the last year of each training window. This filter and estimation scheme results in 722 valid root words in the first estimation window (from 1999 to 2000), and 2,476 valid root words for the last estimation window (from 1999 to 2015).

Finally, we aggregate the tone of each article into a text-based factor F_t , whose weight in (6) is proportional to the market capitalization of common shares outstanding for firms for which a tone observation is available at time t . Our motivation for this choice is twofold. First, as observed in our sample, large firms (based on market capitalization) tend to have more news articles written about

¹⁴Engelberg (2008) finds a similar relationship for the Dow Jones newswire and posits several reasons for this, such as more expansive analyst coverage, catering to client demand (i.e., a larger shareholder base), or even market capitalization internal thresholds for news coverage.

them than small firms. As such, large firms should have a bigger impact on the overall market tone. Second, a market capitalization-weighted text-based normal tone factor mimics the traditional market model used in the event study literature for measuring abnormal returns.¹⁵

Table 1: Number of articles and coverage.

| <i>Panel A. All sources</i> | | | | | | | |
|----------------------------------|-----------------|---------------|-----|--------|---------------|------|--------|
| Bucket | Number of firms | News articles | | | News coverage | | |
| | | Average | Min | Max | Average | Min | Max |
| #1 | 150 | 948 | 4 | 12,496 | 20.67 | 0.76 | 85.26 |
| #2 | 149 | 1,447 | 2 | 11,096 | 28.23 | 0.27 | 81.04 |
| #3 | 149 | 2,645 | 23 | 19,716 | 35.65 | 1.65 | 90.97 |
| #4 | 150 | 10,449 | 20 | 86,560 | 54.82 | 4.32 | 100.00 |
| All | 598 | 3,878 | 2 | 86,560 | 34.85 | 0.27 | 100.00 |
| <i>Panel B. Newswires</i> | | | | | | | |
| Bucket | Number of firms | News articles | | | News coverage | | |
| | | Average | Min | Max | Average | Min | Max |
| #1 | 150 | 680 | 1 | 7,415 | 17.04 | 0.32 | 71.40 |
| #2 | 149 | 1,035 | 2 | 8,450 | 22.84 | 0.27 | 73.55 |
| #3 | 149 | 1,817 | 7 | 12,325 | 28.45 | 1.32 | 80.69 |
| #4 | 150 | 5,715 | 17 | 45,072 | 45.75 | 3.02 | 98.98 |
| All | 598 | 2,315 | 1 | 45,072 | 28.53 | 0.27 | 98.98 |
| <i>Panel C. Newspapers</i> | | | | | | | |
| Bucket | Number of firms | News articles | | | News coverage | | |
| | | Average | Min | Max | Average | Min | Max |
| #1 | 150 | 163 | 0 | 2,427 | 4.81 | 0.00 | 46.63 |
| #2 | 149 | 211 | 0 | 2,256 | 5.68 | 0.00 | 56.79 |
| #3 | 149 | 477 | 0 | 5,044 | 9.53 | 0.00 | 58.53 |
| #4 | 150 | 2,874 | 3 | 34,459 | 26.91 | 0.35 | 99.87 |
| All | 598 | 933 | 0 | 34,459 | 11.75 | 0.00 | 99.87 |
| <i>Panel D. Web publications</i> | | | | | | | |
| Bucket | Number of firms | News articles | | | News coverage | | |
| | | Average | Min | Max | Average | Min | Max |
| #1 | 150 | 104 | 0 | 2,654 | 3.14 | 0.00 | 42.81 |
| #2 | 149 | 199 | 0 | 2,633 | 5.84 | 0.00 | 40.43 |
| #3 | 149 | 351 | 0 | 2,347 | 8.40 | 0.00 | 58.97 |
| #4 | 150 | 1,858 | 0 | 31,024 | 20.18 | 0.00 | 99.87 |
| All | 598 | 629 | 0 | 31,024 | 9.40 | 0.00 | 99.87 |

This table reports the average, minimum and maximum number of news articles, as well as the daily percentage of news media coverage, for the 598 non-financial historical constituents of the S&P 500 over the period 2000-Q1 to 2016-Q4. We split the constituents into four buckets, where the 1st (4th) bucket is composed of the firms with the lowest (highest) average market-capitalization value when they were constituents of the index.

3.4 Tone measurement for earnings press releases and earnings calls

We compute the tone of earnings press releases and earnings calls in two ways. First, we use the lexicon by Loughran and McDonald (2011, hereafter LM) as in Huang, Teoh, and Zhang (2013). Second, we use the same GWP model as the one used to compute the tone of media news.¹⁶ The first approach is

¹⁵The results remain qualitatively similar when using an equally-weighted aggregation scheme.

¹⁶A third approach is to use two GWP models: one trained on earnings press releases and one on earnings call transcripts. This methodology has the natural advantage of linking directly the words in the corporate disclosures and call transcripts to the market returns. The pitfall with this approach is the low frequency of the publications leading to a relatively low number of texts that we can use for the model-fitting compared to media news articles. The optimal choice of the tone estimation method is case-specific. In the Internet Appendix, Section D, we report the results when computing the tone using GWP models trained on earnings press releases and earnings calls. We find that, for our corpus, using the GWP on media news leads to higher numerical stability and better out-of-sample performance.

standard to compute the tone of corporate disclosures in accounting and finance. The second is better aligned with our definition of tone as the expected stock return implied by the textual document. The LM and GWP methods applied to earnings press releases and earnings calls allow us to construct four variables to control for the information disclosed by the firm to isolate information provided by the media. We call these variables *EPRLM*, *EPRGWP*, *ECLM*, and *ECGWP*.

4 Media coverage and abnormal tone around earnings announcements

Earnings announcements provide essential information to investors (Basu et al., 2013). The media provide substantial coverage of earnings announcements, which suggests they play a critical role in distributing information about announcements (see Tetlock, Saar-Tsechansky, and Macskassy, 2008). In this section, we aim to understand the media environment around earnings announcements better. First, we analyze media coverage near earnings announcements. Second, we look at how media news articles differ from earnings press releases and earnings calls in terms of the topics covered. Third, we analyze the most influential words for the tone obtained with the GWP model when applied to news articles, earnings press releases, and earnings calls. Finally, we analyze the drivers of CAT at and after an announcement through graphical and panel regression analyses.

Our data set contains 23,490 earnings announcements for 598 firms. However, several conditions must be met for an event to be included in our analysis. First, we require that every event-firm observation has a daily tone measure available on the event date, or one day before or one day after. Second, to mitigate potential issues related to missing observations for the estimation of the normal tone model, we require at least ten daily tone measures in the estimation window. Finally, we also require that we have an estimate of the tone for the earnings press release and the earnings call. In total, these filters reduce the number of events to 6,394.

We consider an estimation window of 30 days with an offset of 5 days (i.e., $L = 30$, $K = 5$; the estimation window is $[t_i - 35, t_i - 6]$). We focus our analyses on four post-estimation windows: (i) the pre-event window (i.e., $[t_i - 5, t_i - 2]$), (ii) the event window (i.e., $[t_i - 1, t_i + 1]$), (iii) the short-term post-event window (i.e., $[t_i + 2, t_i + 5]$), and (iv) the long-term post-event window (i.e., $[t_i + 2, t_i + 20]$). This setup allows us to consider the pre- and post-event dynamics of the abnormal tone, as well as the reaction at the time of the event.¹⁷

4.1 Number of news articles

We first analyze the media's publication of firm-related news articles around earnings announcements. We expect a spike in the number of articles released during the event window for all three types of media sources. We also expect heterogeneity in the extent of the spike across the types of sources: Newswires are the primary source of news about firms, whereas newspapers and web publications can decide whether to report about an event. We thus expect more articles from newswires than from newspapers and web publications.

In Figure 4, we display the daily average number of articles published around the event date for the three media types. For newswires, newspapers, and web publications, the average number of articles during the event window are 2.66, 1.16, and 0.65, respectively, compared to averages on non-event days of 1.32, 0.57, and 0.31, respectively. While this is in line with our expectations, we also note a substantial number of articles released in the short-term post-event window. We posit that some media outlets are late in reporting about earnings announcements (see, e.g., Tetlock, 2011). These articles may contain further analysis or updated information, such as analyst comments on the financial results, on the days following an announcement.

¹⁷The total length of the estimation and event windows is 61 days, which is the typical length of time between a firm's quarterly earnings announcements. This ensures we have non-overlapping events when considering a single firm.




Figure 4: Average number of articles around the event date. This figure shows the (cross-section) average number of articles relative to the event date by source types across 6,394 quarterly earnings announcement events. Day 0 is the event day, and the gray area represents the event window.

4.2 Coverage of topics

We study the incremental information value of the media by analyzing the difference in content and tone between the newspaper articles and the earnings press releases and earnings calls. The topic analysis is arduous because the class of documents can affect word frequencies across topic and topic prevalence. We address that challenge by using the structural topic model (STM) of Roberts, Stewart, and Airoldi (2016) that allows us to incorporate the effects of metadata (here a variable indicating whether the text is classified as a news article, an earnings press release, or an earnings call transcript) into the topic model. Before applying STM, we process our corpus of texts following Hansen, McMahon, and Prat (2018, p.218) to identify two-word and three-word collocations, which are sequences of words that have a specific meaning. We estimate the STM with 30 topics.¹⁸

In Table 2, we report the ten largest (absolute) differences in topic prevalence measured between the corpus of news articles and the earnings press releases (Panel A) and earnings calls (Panel B). We manually assigned a label for each topic based on the most common words or collocations. First, we see that news media are more likely to cover topics “analysts, legal, innovation, and stock research” compared to earnings press releases and earnings calls. The focus on “analysts” suggests that a specific role of media is in providing context to the firm results by comparing them against analyst recommendations. The higher prevalence of topics “legal, innovation, and stock research” points to the information that usually appears in earnings press releases and earnings calls but is highlighted by the media. It is also the case for topic “board” in earnings calls. On the other hand, themes “accounting” and “earnings” are less prevalent in news articles than in earnings press releases. The same holds for “cash/debt management” in the case of earnings calls. These represent three types of information that the media could have picked up, but in practice, are generally not. Topics such as “entertainment, telecommunication, electricity utility, pulp and paper, and consumer goods” are related to sectors. Finally, note that for earnings call transcripts, more prevalent collocations are typically used to report

¹⁸Using the approach of Mimno and Lee (2014), we find that the optimal number of topics is 54. Hansen, McMahon, and Prat (2018) estimate the optimal number of topics at 70, but reduce this to 40 for ease of interpretation. We follow their reasoning.

the discussion of the call (i.e., labeled “ec-related terms”). This topical analysis illustrates that media focus on certain discussion subjects to contextualize, emphasize, or mitigate specific results compared to earnings press releases and earnings calls.

Table 2: Difference in topic prevalence.

| <i>Panel A. News articles vs. earnings press releases</i> | | |
|---|----------------------|---|
| ΔP | Topic | High-probability word or collocation |
| 0.112 [0.108;0.117] | analysts | analysts, chief_executive, consensus_revenue_estimate, consensus_estimate, profit_per_share |
| 0.083 [0.080;0.087] | legal | court, investigation, lawsuit, attorney, bankruptcy |
| 0.059 [0.056;0.063] | innovation | award, technology, leader, developer, manager |
| 0.056 [0.052;0.060] | stock research | investment_research, sentiment, price_target, economist, rating |
| 0.022 [0.019;0.024] | entertainment | show, radio, entertainment, football, music |
| 0.019 [0.016;0.022] | telecommunication | connection, telecommunication, phone, carrier, network |
| -0.013 [-0.017;0.009] | electricity utility | megawatt, electricity, outage, load, grid |
| -0.022 [-0.025;0.020] | pulp and paper | pulp, tissue, paper, wood, raw_material |
| -0.092 [-0.095;-0.089] | accounting | deferred_income, total_current_liabilities, net_cash, amortization_of_goodwill, provision_for_income |
| -0.234 [-0.238;-0.230] | earnings | news_release, operating_performance, foreign_currency_exchange, other_risks, non-gaap_financial_measure |
| <i>Panel B. News articles vs. earnings calls</i> | | |
| ΔP | Topic | High-probability word or collocation |
| 0.119 [0.116;0.123] | analysts | analysts, chief_executive, consensus_revenue_estimate, consensus_estimate, profit_per_share |
| 0.080 [0.077;0.083] | legal | court, investigation, lawsuit, attorney, bankruptcy |
| 0.075 [0.072;0.078] | innovation | award, technology, leader, developer, manager |
| 0.070 [0.066;0.073] | stock research | investment_research, sentiment, price_target, economist, rating |
| 0.032 [0.020;0.036] | earnings | operating_performance, foreign_currency_exchange, other_risks, non-gaap_financial_measure, news_release |
| 0.021 [0.019;0.023] | board | director, shareholder, committee, holder, annual meeting |
| 0.018 [0.016;0.020] | telecommunication | connection, telecommunication, phone, carrier, network |
| -0.021 [-0.024;-0.019] | consumer goods | merchandise, footwear, jewelry, shopping, fashion |
| -0.063 [-0.065;-0.061] | cash/debt management | receivables, liquidity, working_capital, debt_reduction, cap_ex |
| -0.261 [-0.263;-0.258] | ec-related terms | analyst_presentation_operator, transcription, speakers, follow_up, question |

This table reports topics with the largest absolute difference in prevalence between media news articles published near earnings announcements and earnings press releases (Panel A) or earnings calls (Panel B). Column ΔP reports the topic prevalence's difference between news articles and earnings press releases or earnings calls together with its 99% confidence interval in squared brackets. A positive (negative) value indicates that the topic is more (less) prevalent in news articles than earnings press releases or earnings calls. Column “Topic” reports the topic name (manually assigned). The last column reports the five most common words or collocations (sequence of words with a specific meaning) with a high probability of occurring in the topic. Note that stop words (e.g., “and”, “the”, “to”, etc.) have been removed.

In addition to comparing the textual content, it is also interesting to compare the words driving the tone in news articles, earnings press releases, and earnings calls around earnings announcements. We provide such an analysis in the Internet Appendix, Section E. We also summarize the context surrounding the appearance of the important words.

4.3 Average CAT by level of SUE

Our previous analyses reveal that media report substantially about the firm around earnings announcements. Moreover, we find variations in the topics covered, the ranking of words driving the tone, and the context in which they appear, between news articles and the earnings press releases and earnings calls. Despite these differences, we still expect a positive relationship between the tone and the earnings surprise. To test this, we first look at the evolution of the (cross-section) average CAT conditional on different standardized unexpected earnings (SUE) levels. We split the events into five buckets, where the cutoff points for SUE are based on their quintiles.

In Figure 5, we display the evolution of the average CAT around the event. During the pre-event window, there is no pattern suggesting that CAT can help predict SUE. While the media reaction

appears to start the day before an announcement, this effect is due to overnight earnings announcements. However, during the event window, we find a clear relationship between the level of earnings surprise and the CAT: the media reacts more strongly to negative than positive earnings surprises. For the short-term post-event window, we see a small drift across the negative SUE buckets and the most positive SUE bucket. For the long-term post-event window, we do not see any clear pattern across the buckets.




Figure 5: Evolution of the CAT around the event date. This figure shows the evolution of the (cross-section) average CAT for five SUE buckets over the 6,394 quarterly earnings announcements. The buckets are based on the quintiles of SUE from the lowest (#1) to the highest (#5). Day 0 is the event day, and the gray area represents the event window.

4.4 Earnings information in CAT

The previous graphical analysis indicates that the average CAT is related to the level of earnings surprises. We now use a panel regression to study the predictive value of the information available in earnings press releases and earnings calls for the media's abnormal tone about a firm. We focus on three time horizons for the CAT dependent variable: (i) at the event with $CAT(-1,1)$, (ii) the short-term post-event with $CAT(2,5)$, and (iii) the long-term post-event with $CAT(2,20)$ window.¹⁹ For the explanatory variables, we consider several variables that journalists could report on or be affected by. We consider the earnings surprise, SUE ; the negative earnings surprise indicator, $I[SUE < 0]$; and their interaction, $SUE \times I[SUE < 0]$. We also consider variables capturing the tone of the earnings press release and the earnings call, as these are the source material for journalists. We use $EPRLM$, $EPRGWP$, $ECLM$, and $ECGWP$, where EPR refers to the earnings press release, EC to the earnings call, LM to the lexicon-based measure, and GWP to the GWP-based measure. We include CAT and the cumulative abnormal return, CAR , before and on the event date to control for past media reporting and to determine whether the media tends to focus on past market reaction. To compute CAR , we follow MacKinlay (1997) and calculate abnormal returns using a market model, where the market is defined as in Fama and French (1992).²⁰ We include as additional control variables the return on assets, ROA ; the logarithmic book-to-market ratio, $\log(B/M)$; and logarithmic market value, $\log(M)$;

¹⁹For readability, we remove the event-firm index i for the dependent and explanatory variables.

²⁰The market is the return of a value-weighted portfolio of all CRSP firms incorporated in the United States and listed on the NYSE, AMEX or NASDAQ. Data are available from Kenneth R. French's website at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

and firm and year-quarter fixed effects. Finally, we winzorize the variables SUE , ROA , $\log(M)$, and $\log(B/M)$ at the 1% and 99% levels.

Regression results are reported in Table 3. We find a significant relationship between the media's abnormal tone and the tone in earnings press releases and earnings calls during the event window. The media's abnormal tone tends to be lower (higher) when the tone of the earnings press release or the earnings call is negative (positive). The latter relationship is expected, as the earnings press release and the earnings call are major sources for journalists when writing news articles about the firm's earnings announcement. We also see that the media's abnormal tone tends to be lower when the surprise is negative. This finding is in line with several studies which have identified that the coverage of bad news is disproportionately negative (see, e.g., Soroka, 2006; Garz, 2013; Arango-Kure, Garz, and Rott, 2014). In addition, we find that the media's abnormal tone is positively correlated with its past value. This serial correlation can be attributed to the media's tendency to add context to articles by referring to past information (see Fink and Schudson, 2014; McIntyre, Dahmen, and Abdenour, 2018). We observe a similar media tone momentum for the short- and long-term post-event CAT reported in the second and third columns of Table 3.

Finally, consistent with the view that the media only reports about firms when there is new information, we find that the media's abnormal tone in the long-term post-event window is negatively related to the tone of the earnings press release and the tone of the earnings call. This is in line with the finding of Graffin, Halebian, and Kiley (2016) that firms tend to release unrelated positive news in the days immediately surrounding a potentially negatively received acquisition announcement. They suggest that managers may wish to engage in this sort of reputation management in other contexts that can be easily anticipated, such as negative earnings announcement events.

5 Media tone and stock market reaction around earnings announcements

We now investigate how news articles relate to movements in stock prices and trading volumes. The attribute of interest is the RCAT of media articles, the residual of the panel regression estimated in the previous section. Its construction is analogous to the one of the abnormal tone of earnings press release analyzed in Huang, Teoh, and Zhang (2013) and Arslan-Ayaydin, Boudt, and Thewissen (2016) who use a similar regression to extract the managerial tone inflation in earnings press releases. In media news analysis, the RCAT corresponds to the unexpected media tone and thus reflects the effect of transformation of information by the media rather than dissemination. We denote the RCAT over the event window as $RCAT_{(-1,1)}$.

If the media's transformation of the primary information has some value to investors, $RCAT_{(-1,1)}$ should be positively correlated with the stock price reaction in the three days around an event. Such a positive reaction would be consistent with the price reaction to the (abnormal) tone in earnings press releases (see Henry, 2008; Demers and Vega, 2008; Davis, Piger, and Sedor, 2012; Huang, Teoh, and Zhang, 2013) and earnings calls (see Price et al., 2012). It is also consistent with Engelberg (2008), who observes a positive correlation between stock returns and the tone of articles from newswires at earnings announcements. In the case where $RCAT_{(-1,1)}$ is a proxy for investor sentiment and the behavior of noise traders, a reversal in stock price should be observed (see, e.g., De Long et al., 1990; Tetlock, 2007; Dougal et al., 2012; Garcia, 2013). Huang, Teoh, and Zhang (2013) and Arslan-Ayaydin, Boudt, and Thewissen (2016) observe this pattern on the inflated part of the tone of earnings press releases.

Below we first study the relationship between $RCAT_{(-1,1)}$ and the stock market using a graphical analysis of average CAR based on buckets conditional on $RCAT_{(-1,1)}$. We then use panel regression to test for the incremental information of the media's abnormal tone compared to the quantitative and qualitative information available in the earnings press release and earnings call. We also study whether

Table 3: CAT regression results.

| | <i>CAT</i> (-1,1) | <i>CAT</i> (2,5) | <i>CAT</i> (2,20) |
|--|-----------------------|---------------------|---------------------|
| <i>SUE</i> | 0.192 (0.126) | 0.144 (0.103) | 0.118 (0.243) |
| <i>I</i> [<i>SUE</i> <0] | -0.003*** (0.0005) | 0.00000 (0.0004) | 0.002** (0.001) |
| <i>SUE</i> × <i>I</i> [<i>SUE</i> <0] | -0.362* (0.220) | -0.183 (0.134) | 0.124 (0.286) |
| <i>EPRLM</i> | 0.007 (0.020) | -0.001 (0.021) | -0.055 (0.048) |
| <i>EPRGWP</i> | 0.446*** (0.058) | -0.030 (0.048) | -0.094 (0.127) |
| <i>ECLM</i> | 0.255*** (0.042) | 0.058 (0.045) | -0.002 (0.098) |
| <i>ECGWP</i> | 0.480*** (0.091) | -0.063 (0.083) | -0.361* (0.186) |
| <i>CAT</i> (-5,-2) | 0.223*** (0.026) | 0.203*** (0.023) | 0.639*** (0.054) |
| <i>CAT</i> (-1,1) | | 0.204*** (0.019) | 0.556*** (0.054) |
| <i>CAR</i> (-5,-2) | 0.001 (0.005) | | |
| <i>CAR</i> (-1,1) | | 0.001 (0.002) | 0.003 (0.006) |
| <i>ROA</i> | 0.009 (0.011) | 0.011 (0.010) | 0.047 (0.030) |
| <i>log</i> (<i>B</i> / <i>M</i>) | -0.002* (0.001) | 0.002 (0.001) | 0.003 (0.003) |
| <i>log</i> (<i>M</i>) | 0.0001 (0.0004) | 0.0002 (0.0004) | -0.0004 (0.001) |
| Firm fixed effects | Yes | Yes | Yes |
| Year-quarter fixed effects | Yes | Yes | Yes |
| Observations | 6,394 | 6,394 | 6,394 |
| <i>R</i> ² (×100) | 10.0 | 9.7 | 13.0 |

This table reports the panel regression results of the CAT computed over various time spans on firm- and event-specific variables. Dependent variables are the CAT computed: (i) at the event with *CAT*(-1,1), (ii) short-term post-event with *CAT*(2,5), and (iii) long-term post-event with *CAT*(2,20). Explanatory variables are the earnings surprise, *SUE*; the negative earnings surprise indicator, *I*[*SUE*<0]; the interaction between those two variables, *SUE*×*I*[*SUE*<0]; the tones of earnings press releases and earnings calls, *EPRLM*, *EPRGWP*, *ECLM*, *ECGWP*, where *EPR* refers to earnings press releases, *EC* to earnings calls, *LM* to the lexicon-based measure, and *GWP* to the GWP-based measure; *CAT* and *CAR* measures for various time spans; the return on assets, *ROA*; the logarithm of the book-to-market ratio, *log*(*B*/*M*); the logarithm of the market-capitalization, *log*(*M*); and year-quarter and firm fixed effects. Our sample consists of 8,344 earnings announcements for 598 non-financial S&P 500 constituents from 2000 to 2016. We use a market capitalization text-based normal tone factor model and the market model for the normal return. For each case, we use the estimation window $[t_i - 30, t_i - 6]$, where t_i is the date of event-firm i , to estimate the normal tone and the normal return models. The daily tone is estimated using the GWP methodology, where the variable of interest is the firms' stock returns. Significant explanatory variables are highlighted in gray, and the level of significance is indicated by ** for 5% and *** for 1%. The standard errors are computed using the double-clustered (year-quarter and firm) standard error estimator (Petersen, 2009) and are reported in parentheses below the parameter estimates.

the type of news matters and how *RCAT*(-1,1) relates to the abnormal volume around the earnings announcement. To avoid any look-ahead bias, we estimate the panel regression up to and including the year-quarter in which *CAT*(-1,1) is present to compute *RCAT*(-1,1). Thus, our analysis relies on data from 2004–Q1 to 2016–Q4.

5.1 Average CAR by level of RCAT

We first tackle this question from an exploratory data analysis viewpoint and investigate the evolution of the (cross-section) average CAR conditional on the level of *RCAT*(-1,1) and the direction of the earnings surprise. Specifically, we build six buckets based on *RCAT*(-1,1) terciles conditional on the sign of the earnings surprise (three for positive earnings surprises and three for negative).

In Figure 6, we display the evolution of the average CAR around the event. First, even though $RCAT(-1,1)$ is orthogonal to the information that can be derived from earnings press releases and earnings calls, we find a sizeable positive relationship between the abnormal return and RCAT during the event window. The relationship appears to be stronger when an earnings surprise is negative. Second, in the long-term post-event window, we observe a reversal in stock price when an earnings announcement is negative. This reversal appears to be stronger the lower $RCAT(-1,1)$ is, and it is partial, at least in the 20 trading days after the event. We do not observe a reversal for positive earnings announcements.




Figure 6: Evolution of the CAR around the event date. This figure shows the evolution of the (cross-section) average CAR for six RCAT buckets (three for positive earnings surprises and three for negative earnings surprises) over 4,662 quarterly earnings announcements. The buckets are based on terciles of $RCAT(-1,1)$ from the lowest (#1) to the highest (#3) conditional on the sign of the corresponding value of SUE. Day 0 is the event day, and the gray area represents the event window.

5.2 Stock price reaction information in RCAT

We now perform more extensive analyses using panel regressions to assess the relationship between stock returns and the media's transformation of primary information while controlling for confounding factors. We are interested in both the immediate and delayed responses of the CAR to RCAT. Consistent with our previous analyses, we focus on three time horizons for the CAR dependent variable: (i) at the event with $CAR(-1,1)$, (ii) the short-term post-event window with $CAR(2,5)$, and (iii) long-term post-event window with $CAR(2,20)$. We include the same control variables as in the CAT regressions in Section 4.

We report the results of the regressions in Table 4. First, we find a positive and significant relationship between RCAT and CAR in the event window. An increase of one standard deviation in $RCAT(-1,1)$ is associated with an abnormal return of 0.84% (i.e., the coefficient of 0.825 times the standard deviation of $RCAT(-1,1)$ that is 1.018) at the time of the event. Second, although we do not find that $RCAT(-1,1)$ has a significant relationship with $CAR(2,5)$, we find a negative and significant relationship with $CAR(2,20)$. In this case, an increase of one standard deviation in $RCAT(-1,1)$ is associated with an abnormal return of -0.36% in the next 19 trading days. Thus, we observe a partial reversal of the initial reaction. These results suggest that $RCAT(-1,1)$ conveys information about the market's response around an earnings announcement. This information is incremental to the qualitative and

quantitative information available in the earnings press release and the earnings call, and investors overreact to this information.

Table 4: CAR regression results.

| | <i>CAR</i> (-1,1) | <i>CAR</i> (2,5) | <i>CAR</i> (2,20) |
|--|----------------------|----------------------|----------------------|
| <i>RCAT</i> (-1,1) | 0.825*** (0.109) | -0.029 (0.055) | -0.356*** (0.120) |
| <i>SUE</i> | 5.991*** (0.757) | 0.108 (0.348) | -2.541*** (0.839) |
| <i>I</i> [<i>SUE</i> <0] | -0.019*** (0.003) | 0.001 (0.001) | 0.005 (0.003) |
| <i>SUE</i> × <i>I</i> [<i>SUE</i> <0] | -6.635*** (0.944) | 0.337 (0.708) | 3.219** (1.314) |
| <i>EPRLM</i> | 0.168 (0.109) | 0.010 (0.062) | 0.278* (0.155) |
| <i>EPRGWP</i> | -0.136 (0.249) | 0.102 (0.130) | -0.227 (0.422) |
| <i>ECLM</i> | 0.924*** (0.218) | -0.005 (0.130) | -0.703** (0.319) |
| <i>ECGWP</i> | 1.665*** (0.563) | -0.928*** (0.292) | -0.621 (0.676) |
| <i>CAR</i> (-5,-2) | 0.089 (0.060) | | |
| <i>CAR</i> (-1,1) | | 0.059*** (0.015) | 0.192*** (0.027) |
| <i>ROA</i> | -0.010 (0.086) | 0.141*** (0.051) | 0.296** (0.130) |
| <i>log</i> (<i>B</i> / <i>M</i>) | -0.014* (0.008) | 0.017*** (0.005) | 0.054*** (0.012) |
| <i>log</i> (<i>M</i>) | 0.001 (0.004) | -0.007*** (0.003) | -0.035*** (0.006) |
| Firm fixed effects | Yes | Yes | Yes |
| Year-quarter fixed effects | Yes | Yes | Yes |
| Observations | 4,662 | 4,662 | 4,662 |
| <i>R</i> ² (×100) | 12.5 | 2.8 | 6.8 |

This table reports the panel regression results on the relationship between *RCAT*(-1,1) and the CAR computed over various time spans. Dependent variables are CAR computed: (i) at the event with *CAR*(-1,1), (ii) short-term post-event with *CAR*(2,5), and (iii) long-term post-event with *CAR*(2,20). Missing values of *RCAT*(-1,1) are set to zero. Our sample consists of 23,448 earnings announcements for 598 non-financial S&P 500 constituents from 2000 to 2016, from which 6,617 observations have observable values of *RCAT*(-1,1). See Table 3 for a description of the explanatory variables and details about the estimation.

The above result is consistent with the hypothesis that the abnormal tone attracts the attention of investors, leading to an immediate overreaction and a subsequent price reversal (see Da, Engelberg, and Gao, 2011). The result contradicts the underreaction hypothesis of Engelberg (2008).

5.3 Stock price information in RCAT conditional on media source

Newswires have different business models and target audiences than newspapers and web publications. The former are typically available to professional investors on platforms such as the Bloomberg terminal, whereas newspapers and web publications have a broader reach and are often consulted by retail investors. Hence, it is possible that the relationship between RCAT and CAR is different for newswires than for newspapers and web publications. To test this, we aggregate the $RATC_{d,i,\tau}(-1,1)$ by two groups: one for newswires, and one for newspapers and web publications. We refer to these values as *NWRCAT*(-1,1) for the newswires and *NPRCAT*(-1,1) for newspapers and web publications.

In Table 5, we report results of regressions that separate the effect of RCAT for the different media sources. We find a stock price reaction to RCAT for both groups at the event: an increase of one standard deviation in *RCAT*(-1,1) for newswires is associated with an abnormal return of 0.56% (0.53% for newspapers and web publications). Interestingly, we find a partial reversal for all media sources. This reversal appears to be larger for newswires than for newspapers and web publications,

even though the immediate reaction is economically larger for newspapers. However, at each time horizon, the difference is not significantly different from zero at standard significance levels.²¹

Table 5: CAR regression results for source-specific RCAT.

| | <i>CAR</i> (-1,1) | <i>CAR</i> (2,5) | <i>CAR</i> (2,20) |
|--|----------------------|----------------------|----------------------|
| <i>NWRCAT</i> (-1,1) | 0.785*** (0.133) | -0.076 (0.078) | -0.497** (0.220) |
| <i>NPRCAT</i> (-1,1) | 0.881*** (0.165) | 0.034 (0.078) | -0.168 (0.163) |
| <i>SUE</i> | 5.986*** (0.755) | 0.104 (0.348) | -2.554*** (0.853) |
| <i>I</i> [<i>SUE</i> <0] | -0.019*** (0.003) | 0.001 (0.001) | 0.004 (0.003) |
| <i>SUE</i> × <i>I</i> [<i>SUE</i> <0] | -6.623*** (0.941) | 0.348 (0.706) | 3.252*** (1.218) |
| <i>EPRLM</i> | 0.169 (0.109) | 0.011 (0.062) | 0.281* (0.155) |
| <i>EPRCC</i> | -0.139 (0.249) | 0.099 (0.131) | -0.237 (0.390) |
| <i>ECLM</i> | 0.918*** (0.217) | -0.011 (0.130) | -0.720** (0.348) |
| <i>ECGWP</i> | 1.671*** (0.565) | -0.921*** (0.293) | -0.602 (0.995) |
| <i>CAR</i> (-5,-2) | 0.089 (0.060) | | |
| <i>CAR</i> (-1,1) | | 0.059*** (0.015) | 0.192*** (0.026) |
| <i>ROA</i> | -0.010 (0.085) | 0.142*** (0.051) | 0.298** (0.136) |
| <i>log</i> (<i>B</i> / <i>M</i>) | -0.014* (0.008) | 0.017*** (0.005) | 0.054*** (0.013) |
| <i>log</i> (<i>M</i>) | 0.001 (0.004) | -0.007*** (0.003) | -0.035*** (0.008) |
| Firm fixed effects | Yes | Yes | Yes |
| Year-quarter fixed effects | Yes | Yes | Yes |
| Observations | 4,662 | 4,662 | 4,662 |
| <i>R</i> ² (×100) | 12.5 | 2.8 | 6.9 |

This table reports the panel regression results on the relationship of source-specific *RCAT*(-1,1) on the *CAR* computed over various time spans. Dependent variables are the *CAR* computed: (i) at the event with *CAR*(-1,1), (ii) short-term post-event with *CAR*(2,5), and (iii) long-term post-event with *CAR*(2,20). For the explanatory variables, we consider the contribution of newswires, *NWRCAT*(-1,1); and the contribution of the newspapers and web publications, *NPRCAT*(-1,1). See Table 3 for the description of the other explanatory variables and details about the estimation.

5.4 Abnormal turnover and RCAT

The partial reversal in stock prices points to an overreaction by investors to the incremental information provided by the media. Hou, Xiong, and Peng (2009) argues that one of the necessary conditions for overreaction is attention, which can be proxied by trading volumes (see Barber and Odean, 2008). We thus investigate whether the magnitude of the incremental information conveyed by RCAT (absolute value of RCAT) has any relationship to trading volume around earnings announcements. We use the abnormal logarithmic share turnover, where we compute the abnormal logarithmic share turnover of the event window by using the mean-adjusted model of Campbell and Wasley (1996). Taking the cumulative sum of the abnormal logarithmic share turnovers over relative times leads to the CAST measure.

In Figure 7, we display the evolution of the (cross-sectional) average CAST conditioned on the magnitude of RCAT and the sign of SUE. For all buckets, the CAST increases at the event date. The

²¹The Wald test statistics of equal coefficients is 0.86 (*p*-value is 0.30) for *CAR*(-1,1), 1.33 (*p*-value is 0.18) for *CAR*(2,5), and 0.85 (*p*-value is 0.40) for *CAR*(2,20).

increase is the most pronounced when the magnitude of RCAT is the highest, thus supporting our conjecture that investors overreact to incremental information conveyed by RCAT.




Figure 7: Evolution of the CAST around the event date. This figure shows the evolution of the (cross-section) average CAST for six RCAT buckets (three for positive earnings surprises and three for negative earnings surprises), where the RCAT buckets are based on terciles of $-RCAT(-1,1)$ —from the lowest (#1) to the highest (#3)) over 4,662 quarterly earnings announcement events.

In Table 6, we report results of the regressions for $CAST(-1,1)$. As explanatory variables, we consider the magnitude of $RCAT(-1,1)$ as well as controls used in our previous analyses. We find a significant positive contemporaneous relationship between the magnitude of $RCAT(-1,1)$ and the abnormal volume, again supporting an overreaction. When looking at different media sources, we find that both the incremental information extracted from newswires or newspapers and web publications has a positive and significant relationship with trading volume.²²

6 Conclusion

The news media plays a central role in disseminating information about corporate earnings. While one could try to find anecdotal evidence of news articles that help to explain stock price reactions around earnings announcements, the scientific challenge is to formalize this exercise into a framework that aggregates the information from news articles into statistics that represent the incremental information that can be distilled from information disclosed by the firms, the earnings press releases, and the earnings calls.

We propose a tone-based event study to quantify the added qualitative information provided in news articles about earnings announcements. We apply our novel methodology to S&P 500 firms from 2000 to 2016. We focus on newswires, newspapers, and web publications.

We show that the abnormal tone in media articles helps explain price reactions around an earnings announcement and predicts subsequent price dynamics in the 20 days following the event. While there is a significant positive contemporaneous association between abnormal tone and price movement, the abnormal tone is negatively associated with post-event price movement. We conclude that the

²²The Wald test statistics of equal coefficients is 0.42 (p -value is 0.67).

media provides additional information not included in the information from earnings press releases and earnings calls, and on aggregate, market participants overreact to it. Complementary analyses of the traded volumes support these findings.

Table 6: CAST regression results.

| | <i>CAST</i> (-1,1) | |
|--|----------------------|----------------------|
| — <i>RCAT</i> (-1,1)— | 0.120** (0.054) | |
| — <i>NWRCAT</i> (-1,1)— | | 0.144* (0.079) |
| — <i>NPRCAT</i> (-1,1)— | | 0.196* (0.103) |
| <i>SUE</i> | 0.701** (0.276) | 0.701** (0.275) |
| <i>I</i> [<i>SUE</i> <0] | 0.004*** (0.001) | 0.004*** (0.001) |
| <i>SUE</i> × <i>I</i> [<i>SUE</i> <0] | -1.383*** (0.521) | -1.386*** (0.522) |
| <i>EPRLM</i> | -0.135*** (0.048) | -0.133*** (0.048) |
| <i>EPRGWP</i> | 0.025 (0.151) | 0.026 (0.153) |
| <i>ECLM</i> | 0.010 (0.123) | 0.013 (0.122) |
| <i>ECGWP</i> | -0.730*** (0.243) | -0.733*** (0.241) |
| <i>CAST</i> (-5,-2) | 0.574*** (0.051) | 0.574*** (0.052) |
| — <i>CAR</i> (-5,-2)— | -0.027 (0.021) | -0.027 (0.021) |
| — <i>ROA</i> — | 0.033 (0.040) | 0.034 (0.040) |
| <i>log</i> (<i>B</i> / <i>M</i>) | -0.007** (0.003) | -0.007** (0.003) |
| <i>log</i> (<i>M</i>) | -0.012*** (0.002) | -0.012*** (0.002) |
| Firm fixed effects | Yes | Yes |
| Year-quarter fixed effects | Yes | Yes |
| Observations | 4,662 | 4,662 |
| <i>R</i> ² (×100) | 14.5 | 14.5 |

This table reports the panel regression results on the relationship between the absolute value of *RCAT*(-1,1) (and group equivalent) and *CAST*(-1,1). We use a mean-adjusted model for the abnormal share turnover. See Table 3 for the description of the other explanatory variables and details about the estimation.

References

- Ahmad, K., J. Han, E. Hutson, C. Kearney, and S. Liu. 2016. Media-expressed negative tone and firm-level stock returns. *Journal of Corporate Finance* 37:152–72.
- Algaba, A., D. Ardia, K. Bluteau, S. Borms, and K. Boudt. 2020. Econometrics meets sentiment: An overview of methodology and applications. *Journal of Economic Surveys* 34:512–47.
- Allee, K. D., and M. D. DeAngelis. 2015. The structure of voluntary disclosure narratives: Evidence from tone dispersion. *Journal of Accounting Research* 53:241–74.
- Arango-Kure, M., M. Garz, and A. Rott. 2014. Bad news sells: The demand for news magazines and the tone of their covers. *Journal of Media Economics* 27:199–214.
- Arslan-Ayaydin, Ö., K. Boudt, and J. Thewissen. 2016. Managers set the tone: Equity incentives and the tone of earnings press releases. *Journal of Banking & Finance* 72:132–47.
- Asay, H. S., R. Libby, and K. Rennekamp. 2018. Firm performance, reporting goals, and language choices in narrative disclosures. *Journal of Accounting and Economics* 65:380–98.

- Barber, B. M., and T. Odean. 2008. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies* 21:785–818.
- Basu, S., T. X. Duong, S. Markov, and E.-J. Tan. 2013. How important are earnings announcements as an information source? *European Accounting Review* 22:221–56.
- Borochin, P. A., J. E. Cicon, R. J. DeLisle, and S. M. Price. 2018. The effects of conference call tones on market perceptions of value uncertainty. *Journal of Financial Markets* 40:75–91.
- Campbell, C. J., and C. E. Wasley. 1996. Measuring abnormal daily trading volume for samples of NYSE/ASE and NASDAQ securities using parametric and nonparametric test statistics. *Review of Quantitative Finance and Accounting* 6:309–26.
- Da, Z., J. Engelberg, and P. Gao. 2011. In search of attention. *Journal of Finance* 66:1461–99.
- Davis, A. K., J. M. Piger, and L. M. Sedor. 2012. Beyond the numbers: Measuring the information content of earnings press release language. *Contemporary Accounting Research* 29:845–68.
- De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann. 1990. Noise trader risk in financial markets. *Journal of Political Economy* 98:703–38.
- Demers, E., and C. Vega. 2008. Soft information in earnings announcements: News or noise? Working paper.
- Dougal, C., J. Engelberg, D. Garcia, and C. A. Parsons. 2012. Journalists and the stock market. *Review of Financial Studies* 25:639–79.
- Engelberg, J. 2008. Costly information processing: Evidence from earnings announcements. Working paper.
- Engelberg, J., R. D. McLean, and J. Pontiff. 2018. Anomalies and news. *Journal of Finance* 73:1971–2001.
- Fama, E. F., and K. R. French. 1992. The cross-section of expected stock returns. *Journal of Finance* 47:427–65.
- Fang, L., and J. Peress. 2009. Media coverage and the cross-section of stock returns. *Journal of Finance* 64:2023–52.
- Fink, K., and M. Schudson. 2014. The rise of contextual journalism, 1950s–2000s. *Journalism* 15:3–20.
- Garcia, D. 2013. Sentiment during recessions. *Journal of Finance* 68:1267–300.
- Garz, M. 2013. Unemployment expectations, excessive pessimism, and news coverage. *Journal of Economic Psychology* 34:156–68.
- Graffin, S. D., J. Halebian, and J. T. Kiley. 2016. Ready, AIM, acquire: Impression offsetting and acquisitions. *Academy of Management Journal* 59:232–52.
- Hansen, S., M. McMahon, and A. Prat. 2018. Transparency and deliberation within the FOMC: A computational linguistics approach. *Quarterly Journal of Economics* 133:801–70.
- Heerschop, B., P. van Iterson, A. Hogenboom, F. Frasinca, and U. Kaymak. 2011. Analyzing sentiment in a large set of web data while accounting for negation. In *Advances in Intelligent Web Mastering – 3*, 195–205. Springer Berlin Heidelberg. doi:10.1007/978-3-642-18029-3_20.
- Henry, E. 2008. Are investors influenced by how earnings press releases are written? *Journal of Business Communication* 45:363–407.
- Hou, K., W. Xiong, and L. Peng. 2009. A tale of two anomalies: The implications of investor attention for price and earnings momentum. Working paper.
- Huang, X., S. H. Teoh, and Y. Zhang. 2013. Tone management. *Accounting Review* 89:1083–113.
- Jegadeesh, N., and D. Wu. 2013. Word power: A new approach for content analysis. *Journal of Financial Economics* 110:712–29.
- Jones, M. J., and P. A. Shoemaker. 1994. Accounting narratives: A review of empirical studies of content and readability. *Journal of Accounting Literature* 13:142–84.
- Ke, Z. T., B. T. Kelly, and D. Xiu. 2019. Predicting returns with text data. Working paper.
- Kelley, E. K., and P. C. Tetlock. 2013. How wise are crowds? Insights from retail orders and stock returns. *Journal of Finance* 68:1229–65.
- Kim, S. 2019. The process model of corporate social responsibility (CSR) communication: CSR communication and its relationship with consumers' CSR knowledge, trust, and corporate reputation perception. *Journal of Business Ethics* 154:1143–59.
- Larsen, V. H., L. A. Thorsrud, and J. Zhulanova. 2021. News-driven inflation expectations and information rigidities. *Journal of Monetary Economics* 117:507–20.
- Li, F. 2008. Annual report readability, current earnings, and earnings persistence. *Journal of Accounting and Economics* 45:221–47.

- Livnat, J., and R. R. Mendenhall. 2006. Comparing the post-earnings announcement drift for surprises calculated from analyst and time series forecasts. *Journal of Accounting Research* 44:177–205.
- Loughran, T., and B. McDonald. 2011. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance* 66:35–65.
- MacKinlay, A. C. 1997. Event studies in economics and finance. *Journal of Economic Literature* 35:13–39.
- McIntyre, K., N. S. Dahmen, and J. Abdenour. 2018. The contextualist function: US newspaper journalists value social responsibility. *Journalism* 19:1657–75.
- Merkley, K. J. 2014. Narrative disclosure and earnings performance: Evidence from R&D disclosures. *Accounting Review* 89:725–57.
- Mimno, D., and M. Lee. 2014. Low-dimensional embeddings for interpretable anchor-based topic inference. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1319–28.
- Nimark, K. P., and S. Pitschner. 2019. News media and delegated information choice. *Journal of Economic Theory* 181:160–96.
- Petersen, M. A. 2009. Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies* 22:435–80.
- Porter, M. F. 1980. An algorithm for suffix stripping. *Program* 14:130–7.
- Price, S. M., J. S. Doran, D. R. Peterson, and B. A. Bliss. 2012. Earnings conference calls and stock returns: The incremental informativeness of textual tone. *Journal of Banking & Finance* 36:992–1011.
- Roberts, M. E., B. M. Stewart, and E. M. Airolidi. 2016. A model of text for experimentation in the social sciences. *Journal of the American Statistical Association* 111:988–1003.
- Shapiro, A. H., M. Sudhof, and D. Wilson. 2017. Measuring news sentiment. Working paper.
- Shiller, R. J. 2017. Narrative economics. *American Economic Review* 107:967–1004.
- Soroka, S. N. 2006. Good news and bad news: Asymmetric responses to economic information. *Journal of Politics* 68:372–85.
- Stone, P. J., and E. B. Hunt. 1963. A computer approach to content analysis: Studies using the general inquirer system. In *Proceedings of the May 21–23, 1963, Spring Joint Computer Conference*, 241–56. ACM.
- Tetlock, P. C. 2007. Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance* 62:1139–68.
- . 2010. Does public financial news resolve asymmetric information? *Review of Financial Studies* 23:3520–57.
- . 2011. All the news that’s fit to reprint: Do investors react to stale information? *Review of Financial Studies* 24:1481–512.
- Tetlock, P. C., M. Saar-Tsechansky, and S. Macskassy. 2008. More than words: Quantifying language to measure firms’ fundamentals. *Journal of Finance* 63:1437–67.
- Wang, J., H. T. Shen, J. Song, and J. Ji. 2014. Hashing for similarity search: A survey. Working paper.