Received: 12 December 2018 | Accepted: 4 September 2024

DOI: 10.1111/1911-3846.12986

ARTICLE

CAAA ACPC CONTEMPORARY ACCOUNTING RESEARCH | RECHERCHE COMPTABLE CONTEMPORAINE

# Discretionary dissemination on Twitter

**Richard M. Crowley**[1] | **Wenli Huang**[2] | **Hai Lu**[3,4]

[1]School of Accountancy, Singapore Management University, Singapore, Singapore

[2]College of Professional and Continuing Education, Hong Kong Polytechnic University, Hung Hom, Hong Kong

[3]Rotman School of Management, University of Toronto, Toronto, Ontario, Canada

[4]Guanghua School of Management, Peking University, Beijing, China

**Correspondence**
Hai Lu, University of Toronto, 105 St. George Street, Toronto, ON M5S 3E6, Canada.
Email: hai.lu@rotman.utoronto.ca

## Abstract

The study provides large-scale descriptive evidence on the timing and nature of corporate financial tweeting. Using an unsupervised machine learning approach to analyze 24 million tweets posted by S&P 1500 firms from 2012 to 2020, we find that firms are more likely to tweet financial information around significantly negative or positive news events, such as earnings announcements and the filing of financial statements. This convex U-shaped relation between the likelihood of posting financial tweets and the materiality of accounting events becomes stronger over time. Whereas research based on early samples concludes that firms are less likely to disseminate financial information on Twitter when the news is bad and material, the symmetric dissemination behavior we find suggests that these conclusions should be revised. We also show that a machine learning algorithm (Twitter-Latent Dirichlet Allocation) is superior to a dictionary approach in classifying short messages like tweets.

**KEYWORDS**
disclosures, discretionary dissemination, social media, Twitter

## La diffusion sélective sur Twitter

## Résumé

Cette étude présente une quantité importante de données descriptives sur la fréquence et le contenu des gazouillis financiers des entreprises. Utilisant une approche d'apprentissage automatique non supervisée, les auteurs analysent 24 millions de gazouillis publiés par les sociétés du S&P 1500 entre 2012 et 2020. Ils constatent que les entreprises ont tendance à diffuser des informations

Accepted by Eddie Riedl.

financières lors d'évènements marquants négatifs ou positifs, comme les annonces de résultats ou la publication des états financiers. Cette relation convexe en forme de U entre la probabilité de publier des gazouillis financiers et l'importance des évènements comptables s'intensifie avec le temps. Alors que les recherches réalisées sur des échantillons initiaux démontrent que les entreprises ont moins tendance à diffuser des informations financières sur Twitter en cas de nouvelles importantes et négatives, le comportement symétrique observé dans cette étude suggère que ces résultats doivent être réexaminés. De plus, les auteurs démontrent que l'algorithme d'apprentissage automatique (Twitter-LDA) est plus efficace qu'une approche basée sur un dictionnaire pour classifier les minimessages comme les gazouillis.

**MOTS-CLÉS**

communications d'informations, diffusion sélective, médias sociaux, Twitter

**JEL CLASSIFICATION**

G14, L30, M14, M15, M40

## 1 | INTRODUCTION

News on firm fundamentals is an important information source for investors. However, information disclosed by firms does not always reach the public efficiently. As a result, firms have incentives to continue disseminating disclosed information via various channels. Over the past decade, social media platforms have become so popular that most public firms have created social media accounts and use them to interact with their followers. One platform, Twitter, has a few unique features that make it attractive to firms.[1] Twitter allows only short and concise messages, along with embedded hyperlinks, images, and videos. It enables firms to initiate direct communication with a large network of followers. The platform was created in October 2006 and quickly gained momentum (Twitter, 2011), and the wide adoption of Twitter by firms has motivated researchers to examine the consequences of disseminating information on the platform (e.g., Blankespoor et al., 2014; Jung et al., 2018; Lee et al., 2015). While prior research has explored firms' strategic information disclosure across various channels, there is insufficient understanding of firms' dissemination of financial information on social media, such as Twitter. In the context of our study, we define financial information as information related to firm fundamentals or financial outcomes (including outcomes in financial markets). Our study investigates whether firms tweet financial information in conjunction with mandatory accounting events, and whether firms are equally likely to disseminate extreme good and bad news.

Managers make strategic decisions on the information content and channels of disclosure to mitigate information asymmetry between managers and investors or to better serve managers' own interests. One of these strategic decisions is whether managers disclose good news and delay, suppress, or accelerate bad news. Analytical and empirical studies on this phenomenon have produced mixed results (see Beyer et al., 2010, for a comprehensive review). While firms

[1]At the time of writing, Twitter is now known as X. As the data of our study cover only the period in which the platform was known as Twitter, we refer to it as Twitter throughout the paper.

may have similar incentives for information dissemination and disclosure (Jung et al., 2018), the two concepts are distinct. It is possible that firms have more incentive to further disseminate good news than bad news, just as they have more incentive to disclose good news and delay bad news. Alternatively, firms may have incentives to avoid both under- and over-valuation; thus, they may disseminate both good and bad news. Ex ante, it is unclear which incentives prevail when firms disseminate financial information on Twitter, leaving this an important empirical question.[2]

Prior literature provides mixed evidence. Studying tweets by 85 technology firms from March 2007 through September 2009, Blankespoor et al. (2014) find that tweets with hyperlinks to firm press releases on earnings are not associated with the direction or magnitude of the news. In contrast, using a sample of 406 firms covering Q1 2010 through Q1 2013, Jung et al. (2018) show that managers are less likely to disseminate earnings information on Twitter when earnings news is bad and when the magnitude of the bad news is higher. The difference in the two studies' results could be due to differences in sample composition, sample period, and classification of tweets. For instance, firms may have exhibited limited strategic behavior in early years when the use of social media for dissemination was not well understood by them. Subsequently, as more firms became familiar with dissemination via social media, strategic behavior may have emerged. Furthermore, firms' tweeting activities may have evolved in response to the rapid increase in the number of users. We exploit the availability of tweet data from a longer and later sample period, which allows us to observe the evolution of strategic dissemination behavior over time. To enable better identification of financial information dissemination on Twitter, we introduce a newly developed machine learning (ML) approach to the literature.

Using a complete sample of 24 million tweets posted by S&P 1500 firms with active Twitter accounts from 2012 to 2020, we investigate (1) whether firms post financial information on Twitter, which we refer to as financial tweets, in conjunction with earnings announcements and annual and quarterly filings (henceforth 10-K/10-Q filings) and (2) whether posting financial tweets depends on the materiality of the events irrespective of whether the events involve bad or good news.[3] To enhance our conclusions, we further explore the intraday pattern of financial tweets around these events. Our investigation provides large-sample evidence consistent with firms strategically disseminating financial information on Twitter in ways that differ from the patterns identified in prior studies.

Our results show that firms have an increased likelihood of disseminating financial tweets around earnings announcements and 10-K/10-Q filings. The likelihood of posting financial tweets and the materiality of news follows a symmetric U-shaped relation: firms are more likely to tweet on days with either extreme good or extreme bad news than on days when the news is relatively moderate. We further explore the time trend of the U-shaped relation with regression analysis for each year from 2012 to 2020. We find that firms react to both extreme good and bad news in earlier years of our sample period, but the U-shaped pattern is more pronounced in later years. When we compare the pattern for the firms that joined Twitter before 2012 and the firms that created Twitter accounts during our sample period, we also find that the results are significantly stronger for the late adopters. This finding suggests that the strategic dissemination behavior evolves over time. In addition, our intraday analyses show that firms are more likely to tweet in the several hours after earnings announcements, but they tweet in the hours both before and after 10-K/10-Q filings. This suggests the presence of strategic dissemination in shorter windows around material events.

[2]On April 2, 2013, the SEC issued guidance confirming that firms can use social media to announce new information in compliance with Regulation Fair Disclosure (SEC, 2013). Although tweets may consist of both disclosure of new information and dissemination of news already disclosed elsewhere, the brevity of messages on Twitter is more consistent with firms summarizing or using hyperlinks to disseminate existing information. That over 80% of firms' financial tweets contain hyperlinks suggests that dissemination is a prevalent use case.

[3]We use the term "materiality" to denote the magnitude of information content as measured by the market responses (positive or negative) to the news events or by the magnitude of earnings surprise (positive or negative).

To identify the nature of each tweet, we use an ML algorithm (Twitter-Latent Dirichlet Allocation [LDA]) to classify each tweet into one of the following categories: financial, non-financial business, and other. We focus primarily on tweets containing financial information. The Twitter-LDA approach is an unsupervised Bayesian algorithm used to learn a set of topics discussed across a collection of tweets. It is suitable for big data analysis due to its self-learning nature and works particularly well for our study, which examines millions of short messages. We compare this ML approach against a dictionary approach, as dictionaries are commonly applied in large-sample studies. After a manual validation, we find that Twitter-LDA classifies tweets much more accurately than the dictionary approach.

Our study makes three major contributions. First, we contribute to the information dissemination literature. Using a large recent sample and a refined methodology, we find that firms disseminate symmetrically around extreme good and bad news from material corporate events, using financial tweets as a proxy of financial information dissemination. This new finding extends our understanding from existing literature that firms strategically avoid disseminating bad news around major corporate events (Jung et al., 2018).

Second, our study provides large-scale descriptive evidence on the timing and nature of firms' dissemination of information on Twitter. Our study complements the results of Blankespoor et al. (2014), who find no evidence of strategic behavior in the dissemination of financial information on Twitter, and the results of Jung et al. (2018), who find that firms avoid tweeting in quarters when the magnitude of negative earnings surprise is significant. Our study also shows that firms' dissemination behavior may evolve over time. A broader implication of our finding is that firms' strategic behavior and the information content of firms' tweets may change as social media platforms become mature and attract more followers.

Third, our study is the first to use Twitter-LDA to process a large volume of tweets to identify financial and non-financial topics. We highlight the differences between an ML and dictionary approach in processing social media data. The ML approach improves classification accuracy while offering a more researcher bias–free assessment of the content of the tweets. Using a more accurate classification helps us refine our beliefs on how firms strategically disseminate financial information on social media. Furthermore, this method offers researchers an automated approach to examine a wide variety of content on Twitter.

## 2 | LITERATURE AND HYPOTHESIS DEVELOPMENT

The voluntary disclosure literature provides mixed predictions on whether managers behave strategically when disclosing good and bad news. Assuming that managers maximize stock price, early analytical studies suggest that firms have disclosure thresholds and managers tend to disclose good news and suppress bad news (Dye, 1985; Verrecchia, 1983). However, assuming that managers' incentives are aligned with those of investors (e.g., that both groups wish to avoid over- and under-valuation), Hummel et al. (2021) show that managers disclose both extreme good and extreme bad news.

These contradictory views are documented in the empirical disclosure literature. In general, managers are portrayed as strategic and their disclosure decisions as discretionary (Beyer et al., 2010; Fields et al., 2001; Healy & Palepu, 2001; Leuz & Wysocki, 2016). Several studies on conference calls examine management's strategic communication and its association with information content (Hollander et al., 2010), firm performance (Mayew & Venkatachalam, 2012), and financial fraud and misreporting (Hobson et al., 2012; Larcker & Zakolyukina, 2012). Other studies examine the strategic decision to disclose good and bad news. Kothari et al. (2009) show that managers delay releasing bad news, while other studies provide evidence that managers disclose material bad news promptly due to litigation concerns (e.g., Skinner, 1994, 1997). In the setting of voluntary earnings guidance, Anilowski et al. (2007) show that managers are more likely to issue downward guidance when firms are performing poorly.

When it comes to financial information dissemination, empirical and theoretical evidence is more limited. There are no prior analytical models predicting managers' strategic behaviors. However, one can argue that managers' incentives for dissemination largely overlap with their incentives for disclosure (Jung et al., 2018). Prior studies demonstrate that investors react to financial information dissemination in ways consistent with how they react to financial information disclosure. For instance, Bushee et al. (2003) show that managers' incentives to disseminate are driven, in part, by catering to the information needs of investors, while Rogers et al. (2016) show that investors react to media dissemination of publicly available information. Therefore, we may expect patterns of strategic dissemination to be similar to patterns of disclosure, such as asymmetric disclosure involving the withholding of bad news. In contrast, Blankespoor and deHaan (2020) show that firms' strategic promotion of their CEO can increase dissemination by traditional media and encourage more favorable coverage on days with either the most positive and negative abnormal returns, which suggests that firms have incentives to increase dissemination symmetrically for material news.

Voluntary dissemination on Twitter in conjunction with mandatory disclosure both restricts and benefits firms. While managers have discretion to decide whether to tweet to provide a warning signal and interpretive information, they do not have full control of the timing, since the disclosure is mandatory. Meanwhile, some unique features of Twitter provide firms an effective means of conveying information to investors. For instance, firms can bypass the 140-character constraint, or 280-character constraint starting in November 2017, by embedding hyperlinks in their tweets. In addition, Twitter enables firms to initiate communication with followers directly, significantly reducing investor information search and processing costs. Twitter also allows firms to monitor investors' reactions to tweets, such as likes, retweets, and replies, so that they can revise their dissemination strategies accordingly.

Tweeting offers firms the opportunity to achieve many objectives, such as highlighting important news, promoting the firm and new products, creating a positive social image, maintaining a transparent information environment, and increasing firm visibility by attracting more followers (e.g., Kumar et al., 2013; Toubia & Stephen, 2013). Therefore, firms are likely to be proactive and tweet more frequently when they would like certain information to reach potential investors or customers. For example, firms tweet to mitigate information asymmetry and increase liquidity (Blankespoor et al., 2014), and they tweet to attenuate the negative price reaction to recall announcements (Lee et al., 2015). Firms may also tweet in conjunction with the materiality of the information, as this can allows them to control the conversation surrounding the information, leading to potentially more favorable reaction to the information by stakeholders. This would suggest a symmetric relationship between materiality and dissemination. Furthermore, tweets by individuals reveal associations between investors' attention and firm-level fundamentals. Individuals' tweets can predict firm fundamentals (Bartov et al., 2018; Tang, 2018), and user activity is associated with earnings surprise sensitivity (Curtis et al., 2016). Given the information contained in users' tweets, firms can reasonably expect that investors on Twitter will be listening to what firms disseminate on the platform.

Regarding whether firms behave strategically on social media, Jung et al. (2018) suggest that firms are less likely to disseminate significantly negative earnings news on Twitter. Their study examines earnings-related tweets around earnings announcements. Such strategic disseminating behaviors are not found by Blankespoor et al. (2014), who use an earlier sample consisting of technology firms and focus on the dissemination of firm press releases via hyperlinks. As Twitter has evolved rapidly, and different types of firms have joined the platform at different points in time, the cause of the inconsistency in findings is unclear. Therefore, it is interesting to investigate whether firms have revised their strategic dissemination behavior, as Twitter has evolved to a mature stage with a significantly larger active user base over the past decade. Ideally, such an investigation would be guided by rigorous theoretical predictions. Other than the visibility argument (Merton, 1987), however, we have few theories built on information dissemination to directly inform our predictions.

Financial information in tweets provides greater benefits to investors who have limited resources or skills to search for financial information about firms in the traditional "pull" information system. These investors are interested in material news regardless of whether it is bad or good. In the context of dissemination, managers may prefer to have influence over the communication of good or bad news when it is material. We thus hypothesize that firms are more likely to tweet financial information around financial events with extreme positive or extreme negative information content. Our hypothesis is stated as follows:

> **Hypothesis.** The likelihood of posting financial information on Twitter increases with the materiality of accounting news events, irrespective of the direction of the news (positive or negative).

Ex ante, the outcome for testing our hypothesis is unclear, as many prior studies on strategic disclosure suggest that firms have various incentives to suppress bad news. If firms adopt a consistent strategy in disseminating financial information on social media, they will continue to disclose good news while withholding bad news. However, our hypothesis implies a U-shaped relation between the likelihood of a firm tweeting financial information on Twitter and the materiality of the news events. This is more likely to be found if managers want to control the conversation surrounding material events. To be consistent with prior studies, we focus on earnings announcements and 10-K/10-Q filings, and we partition them based on their materiality.

# 3 | SAMPLE SELECTION AND MEASUREMENT OF TWEETS

## 3.1 | Sample selection

Our sample consists of all 2,199 public firms that were in the S&P 1500 at any point between January 1, 2012, and December 31, 2018, and our analysis covers all tweets posted by these firms from January 1, 2012, through December 31, 2020. We hand-collect the Twitter handles of these firms.[4] Using these handles, we identify the Twitter ID associated with each account via the Twitter API. While Twitter handles can be changed (e.g., after mergers or rebranding), Twitter IDs are permanent, allowing us to track firms across handle changes. In total, we identified Twitter accounts for 1,641 firms: 459 firms that adopted Twitter during our sample period and 1,182 firms that had an account throughout. After we removed 29 firms that made their tweets available only to followers, 99 firms that never tweeted, and 153 firms that do not have sufficient data for all control variables, our data set contains 1,360 Twitter accounts.

To obtain firms' tweets, we used the Twitter API to download all publicly available tweets associated with each Twitter ID. Public access is limited to the 3,200 most recent tweets per account. There were 614 accounts that posted more than 3,200 tweets before we began scraping data in September 2016; we purchased a complete set of tweets for these accounts from GNIP. We combine these two data sources for our analyses of tweets.[5]

[4]A team of research assistants determined firms' Twitter handles using firms' websites, Twitter search, and Google search. The initial collection took place in 2016; an updated collection took place in 2019. Two research assistants checked each firm, and an author of this study manually checked each collected account using the account's details (bio, links, and tweets) to ensure that the account indeed represented the firm.

[5]GNIP was acquired by Twitter in 2014. It was later rolled into Twitter as part of its data vending services. We manually cross-check the tweets downloaded from the API against the tweets from GNIP for a few firms. We find that only advertisements included in the GNIP data are not captured by the API data. This suggests that the data downloaded by the API should be as reliable as those from GNIP for our study.

Our financial data come from six sources. Financial statement, stock, and executive data are from Compustat Fundamentals Annual, CRSP, and ExecuComp, respectively. Earnings announcement dates and times come from Compustat Fundamentals Quarterly and I/B/E/S, respectively. Analyst counts are derived from I/B/E/S. Release dates and times of annual reports (10-K), quarterly reports (10-Q), and 8-K filings are extracted from WRDS SEC Analytics Suite.

We require all observations to have tweeted at least once before or on the given day and have complete information for all control variables. After imposing these restrictions, our final sample contains over 24.6 million tweets across 2.28 million firm-trading days.

## 3.2 | Tweet measures

All tweet measures are calculated at the firm-day level. Daily measures are calculated based on trading days with a 4:00 p.m. cutoff in the Eastern Time Zone. Our main tests focus on *Financial Tweets*, a firm-day-level indicator equal to one when the firm issued at least one tweet in which the text is primarily financial in nature, and zero otherwise. To determine whether the text is primarily financial in nature, we use two methods: an ML algorithm that examines the content of each tweet (indicated as *ML* in analyses), and a dictionary method (indicated as *Dict* in analyses).

The ML algorithm we use for tweet categorization is the Twitter-LDA algorithm of Zhao et al. (2011). This algorithm is based on the LDA algorithm of Blei et al. (2003), which has recently been adopted by several accounting studies (Bao & Datta, 2014; Brown et al., 2020; Crowley, 2016; Hoberg & Lewis, 2017). The LDA algorithm categorizes the thematic content, or topics, within documents in an automated, researcher bias–free manner. LDA determines topics in an unsupervised Bayesian manner by examining how frequently words within documents covary with one another, but it requires the researcher to determine the number of topics for the algorithm to find. Twitter-LDA extends the basic LDA model to work with shorter "documents" in the form of tweets by incorporating correlations between words within and across Twitter users. This twist on the standard LDA implementation is important because the brevity of tweets means that very few words covary at the document level. Taking user-level word covariance into account dramatically increases the model's understanding of which words fit together, and thus it can provide human-understandable topic groupings for tweet-level data.

We run the Twitter-LDA algorithm to detect 60 topics among the firms' tweets.[6] We then manually classify the topics, identifying one topic that discusses financial information, 42 topics related to non-financial business information such as marketing activities, and 17 other topics. For each topic, Twitter-LDA computes a weighted dictionary that we use to assign each tweet to a topic. We classify a tweet as coming from the topic for which it has the highest weight. As our primary focus is on financial tweets, our analysis focuses on tweets matching the financial topic. In Appendix 2, we provide the top 20 words from each subcategory of tweets from our classification and the text of nine sample tweets from our financial topic. Among the extreme positive financial tweets presented, the first two use simple descriptive language, saying that earnings have been announced and linking to the press release or directly stating financial results. The third tweet includes "bragging," mentioning that the company is "on track to generate industry leading cash flow growth." Among the extreme negative tweets, the first again uses simple descriptive language. The next two messages offer a justification for the negative

[6]We chose 60 topics by running models with varying numbers of topics from 50 to 100, optimizing for the clarity of the financial topic. The optimal number of topics in this process was 60, which provides one clear financial topic.

outcome, suggesting promising aspects of the business going forward.[7] Appendix S1 presents example tweets from other subcategories of tweets (e.g., marketing tweets).[8]

We also present our results using a dictionary measure of financial tweeting that is constructed from two word lists from Jung et al. (2018): a list of earning-related terms and a list of investor-focused terms. We flag a tweet as financial if it contains any of the eight earnings-related terms or if it contains any of 11 investor-focused terms.[9] For robustness, we present results using only the earnings-related terms in Tables S1–S5 in Appendix S1. Three of the earnings-related words included in the dictionary measure—"earnings," "quarter," and "results"—are also in the top 20 words in the financial topic from the Twitter-LDA algorithm. The other words have lower weightings, either due to infrequent usage or lack of discriminating power across topics. One word, "customer," is instead associated with customer support.

Dictionary approaches excel when two conditions are met: (1) select words, *n*-grams, or phrases can accurately pick up the majority of text related to the construct of interest (minimizing Type II error) and (2) text unrelated to the construct of interest generally does not include those words (minimizing Type I error). In contrast, ML approaches excel when contexts are not easy to distinguish, or when associations across words are needed to accurately classify the text. While ML approaches may be noisier than human classification, on complex tasks where dictionaries cannot capture human judgment well, ML generally outperforms dictionaries.

An example of a context that is well suited to a dictionary approach is presented in Hassan et al. (2024), who use it to identify discussion of Brexit in earnings conference calls. As the term "Brexit" is standardized as the primary way to refer to the UK's departure from the European Union, a single word dictionary of "Brexit" acts as a good discriminator: it is sufficient to identify related discussion (low Type II error), and it is unique to discussing Brexit (minimal Type I error). In contrast, our setting has some words that are good discriminators, such as "earnings," and some words that are not, such as "income" (confounded by discussion of financial services), "sales" (confounded by discussion of discounts), and "quarter" (used in many general circumstances). Thus, neither a single word nor a single set of words can discriminate well in our setting. Twitter-LDA, however, does not rely on single words to classify tweets, and it does not require a word to have only one meaning in the corpus. For words that are related to multiple concepts, Twitter-LDA assigns a high weight to all related concepts for the word. When a word has high weights across multiple topics, Twitter-LDA uses the rest of the words in the text to determine which topic is being discussed. For example, Twitter-LDA correctly identifies "<u>sales</u> growth and <u>earnings</u> growth" as financial, while it classifies "shop our <u>quarter</u> end <u>sales</u>" as marketing, even though both phrases include two words from our dictionary approach. For the financial phrase, Twitter-LDA's classification is driven by the words "growth" and "earnings," while its classification of the marketing phrase is driven by the words "sales" and "shop."

To understand the consistency of the *Financial Tweets* measure based on the two different approaches, we present a comparison of the tweets flagged by Twitter-LDA and the dictionary approach in Table 1. Panel A shows the overlap between the measures at the tweet level. Of the 24.6 million tweets in our sample, 110,732 and 999,464 tweets are labeled as financial by our ML and dictionary approaches, respectively. The overlap between the measures is only

[7]In an untabulated analysis, we manually examined 50 randomly selected tweets from each category (extreme negative, non-extreme, and extreme positive) around earnings announcements. We classify each tweet into one of three categories based on its language: (1) simple descriptive language, discussing in a matter-of-fact manner; (2) justification, explaining negative performance or refocusing on positive performance; and (3) bragging, touting firm accomplishments. Across all categories, most tweets use simple descriptive language. Among the extreme negative tweets, 82% use simple descriptive language, 16% offer justifications, and 2% include bragging. Among the extreme positive tweets, 86% use simple descriptive language, 2% offer justifications, and 12% include bragging. Lastly, among the non-extreme tweets, 88% use simple descriptive language, 4% offer justifications, and 8% include bragging.

[8]See the Supporting Information.

[9]The earnings-related terms include the following: earnings, eps, profit, income, revenue, sales, results, or quarter. The investor-focused terms include ceo, ceos, executive, executives, dividend, dividends, board, boards, new product, new products, launch, launched, launching, launches, acquisition, acquisitions, merger, mergers, repurchase, repurchased, repurchasing, investment, investments, customer, and customers.

**TABLE 1** Comparison of financial tweet classifications: ML versus dictionary.

**Panel A: Agreement between measures, tweet level**

| Count (percentage) | *Financial Tweets, Dict* = 0 | *Financial Tweets, Dict* = 1 | Total |
|---|---|---|---|
| *Financial Tweets, ML* = 0 | 23,487,573 | 965,673 | 24,453,246 |
| | (95.6%) | (3.93%) | (99.5%) |
| *Financial Tweets, ML* = 1 | 76,941 | 33,791 | 110,732 |
| | (0.31%) | (0.14%) | (0.45%) |
| Total | 23,564,514 | 999,464 | 24,563,978 |
| | (95.9%) | (4.07%) | (100.0%) |

Measure overlap: 33,791/(965,673 + 76,941 + 33,791) = 3.14%

**Panel B: Agreement between measures, day level**

| Count (percentage) | *Financial Tweets, Dict* = 0 | *Financial Tweets, Dict* = 1 | Total |
|---|---|---|---|
| *Financial Tweets, ML* = 0 | 1,923,494 | 282,562 | 2,206,056 |
| | (84.4%) | (12.4%) | (96.8%) |
| *Financial Tweets, ML* = 1 | 30,583 | 42,091 | 72,674 |
| | (1.30%) | (1.85%) | (3.19%) |
| Total | 1,954,077 | 324,653 | 2,278,730 |
| | (85.8%) | (14.3%) | (100.0%) |

Measure overlap: 42,091/(282,562 + 30,583 + 42,091) = 11.8%

**Panel C: Measure complementarity**

| Count (percentage) | *Earnings Ann* | *Form 10-K, 10-Q* |
|---|---|---|
| *Agree, Both* = 0 | 25,914 | 27,591 |
| | (72.0%) | (84.4%) |
| *Disagree* | 6,566 | 6,988 |
| | (18.3%) | (16.7%) |
| *Agree, Both* = 1 | 3,496 | 1,731 |
| | (9.72%) | (4.92%) |

Measure overlap on earnings announcement days: 3,496/(3,496 + 6,566) = 34.7%
Measure overlap on 10-K or 10-Q filing days: 1,731/(1,731 + 6,988) = 19.9%

**Panel D: Manual classification of financial tweets**

| | Number of tweets in sample | Number of tweets checked | Manual classification: Financial tweets |
|---|---|---|---|
| Tweets only in ML | 76,941 | 200 | 101 |
| | | | (50.5%) |
| Tweets only in dictionary | 965,673 | 200 | 28 |
| | | | (14.0%) |
| Tweets in both measures | 33,791 | 200 | 124 |
| | | | (62.0%) |
| Tweets in neither measure | 23,648,834 | 400 | 0 |
| | | | (0.00%) |

*Note*: Panel A shows a confusion matrix between *Financial Tweets, ML* and *Financial Tweets, Dict* at the individual tweet level, showing both counts of tweets and percentages. Panel B shows a confusion matrix between *Financial Tweets, ML* and *Financial Tweets, Dict* at the firm-day observation level, showing both counts of tweets and percentages. Panel C shows the extent to which various events coincide with varying levels of agreement between the *Financial Tweets, ML* and *Financial Tweets, Dict* measures at the firm-day observation level. Panel D presents the results of a manual check of 1,000 tweets from our sample by an independent reviewer. The independent reviewer was instructed to classify tweets as financial if "the tweet is seemingly targeted at investors and conveys something financial in nature." Variable definitions for all variables are included in Appendix 1. Methodology for *Financial Tweets, ML* is discussed in Section 3.2 and Appendix 2, while methodology for *Financial Tweets, Dict* is discussed in Section 3.2. All figures in parentheses are percentages related to the count right above the figure. Measure overlap represents the number of observations with agreement across the ML and dictionary measures divided by the number of observations with either measure flagging the observation.

33,791 tweets, or 3.14% of the tweets flagged by either approach. The low percentage of overlap between the tweets flagged by both approaches suggests a significant amount of noise in classification. Panel B presents the information from Panel A aggregated at the day level. Again, Twitter-LDA flags fewer observations as having financial tweets than the dictionary approach, though the higher overlap between approaches (11.8% of the tweets flagged by either approach) indicates less noise at the day level. Panel C examines the agreement of the two measures in the context of earnings announcements and 10-K/10-Q filings. On earnings announcement (10-K/10-Q) days the measures tend to agree more (less), with the overlap between approaches at 34.7% (19.9%) of the tweets flagged by either approach.

To further explore the reasons for the variation in tweets flagged by each measure, we conducted a validation exercise, presented in Table 1, Panel D. An independent reviewer read through a sample of 1,000 randomly selected tweets from our sample, stratified into four groups: 200 tweets flagged by both measures, 200 tweets flagged only by Twitter-LDA, 200 tweets flagged only by the dictionary approach, and 400 tweets not flagged by either measure. The best-performing class is the tweets flagged by both measures (62.0% accuracy), followed by the tweets flagged only by Twitter-LDA (50.5% accuracy). Tweets flagged only by the dictionary approach are much less likely to be financial in nature (14.0% accuracy). Weighting by the number of flagged tweets from each class, we can extrapolate out the expected accuracy within flagged tweets (one minus the Type I error rate, i.e., precision) for Twitter-LDA at 54.0% and for the dictionary method at 15.6%. Thus, the ML approach appears much more accurate within the sample of flagged tweets. For tweets not flagged by either measure, we see that none of the 400 checked tweets were financial in nature, which indicates a low Type II error rate outside the measures. The low Type II error rate, paired with the lower number of tweets flagged by the Twitter-LDA measure in Table 1, Panels A and B, and the low accuracy within the tweets flagged only by the dictionary, suggests that the ML approach is much more precise than the dictionary approach.

Lastly, we include some other measures derived from Twitter data for an additional test examining the presence of hyperlinks and as controls. For controls, we include an indicator variable (*Verified*) showing whether a firm has a verified account. Verified accounts have been vetted by Twitter for their authenticity and are “an account of public interest.”[10] We also control for the number of followers (*Followers*) and the number of accounts that the firm follows (*Friends*). The two measures capture the popularity of the Twitter account and are included as the log of one plus each measure. These metadata items are only available at the point in time of data collection. As such, we observe variation in these measures only from 2017 onward. We also construct two other control variables, *Total Tweets*, the total number of tweets the firm posted up to the given day, and *Recent Tweets*, the number of tweets that the firm had posted on Twitter over the prior week (5 trading days). The two variables capture firms’ overall and recent activity on Twitter. *Total Tweets* is included as the log of one plus the measure.

# 4 | RESEARCH DESIGN

For all our primary tests, we construct a daily data set of the measures described in Section 3.2. We first use a logistic regression to examine whether firms post financial information on Twitter around major events (earnings announcements, 10-K/10-Q filings, non-earnings announcement 8-K filings),[11] as given by Equation (1):

[10]It is unclear why some accounts are not verified, though it is possible that the firms did not consider verification a value-added service and did not seek verification. Starting in April 2023, verification was replaced with a paid system. This change does not impact our sample as we acquired all data on verification status before the change was made.

[11]8-K filings are included only in our first test in Table 3. As expected, non-earnings 8-Ks do not increase the likelihood of a firm posting a financial tweet under the ML measure, so we drop them from subsequent analyses.

$$\Phi^{-1}(Financial\ Tweets_{i,d}) = \alpha + \beta_1 Event_{i,d} + \beta_2 Twitter\ Controls_{i,d} + \beta_3 Financial\ Controls_{i,d} + Fixed\ Effects_{i,d}, \quad (1)$$

where $i$ represents firms and $d$ represents trading days. The dependent variable is whether the firm posts a financial tweet on the given trading day (*Financial Tweets*), and the variable of interest is *Event*, measured using indicator variables for earnings announcements, 10-K/10-Q filings, and 8-K filings, respectively.

To test our hypothesis, we use the same data and structure as in Equation (1) but examine the impact of the events on firms' daily tweeting behavior conditional on the materiality of the event (captured by event-day abnormal return), as given by Equation (2):

$$\Phi^{-1}(Financial\ Tweets_{i,d}) = \alpha + \sum_{q \in quintiles} \beta_{1,q} Outcome_{i,d,q} \times Event_{i,d} + \beta_2 Twitter\ Controls_{i,d} + \beta_3 Financial\ Controls_{i,d} + Fixed\ Effects_{i,d}. \quad (2)$$

In Equation (2), the variables of interest are the interaction between three quintile-based indicators (*Outcome*) representing the materiality and direction of news events and the events themselves. To examine the materiality of events, we use market model abnormal returns on the trading day of the event. To determine events that the market perceives as more material, our main tests examine tweeting behavior by quintiles of abnormal return within firm over our sample period. We classify an event into a specific quintile based on the abnormal return on the event day relative to other days for each firm in our sample period. For parsimony, we combine the middle three quintiles into one group. We refer to these groups in our analyses as *Extreme Negative AR*, *Non-extreme AR*, and *Extreme Positive AR*, which represent the lowest, middle three, and highest quintiles, respectively. As the variables of interest are interactions between our events and the materiality measures, we also check marginal effects for all such coefficients.

In both Equations (1) and (2), we use the same controls and fixed effects structure. For Twitter controls, we include the control measures discussed in Section 3.2. Our financial controls follow a standard list of control variables based on Jung et al. (2018): firms' most recently reported size (log of assets, *Size*), ROA (*ROA*), market-to-book ratio (*MB*), debt to assets (*Debt*), revenue growth (*Growth*), advertising expense (*Advertising*), return volatility over the past 3 months (*Return Volatility*), age (*Firm Age*), the number of analysts (log(# *Analysts*)$_{\perp}$), and CEO age (*CEO Age*). All variables are defined in Appendix 1. We include both year and month fixed effects (as Twitter activity rapidly increased during the sample period and exhibits seasonality across months) and industry fixed effects (as firms in some industries, such as IT, are more likely to tweet).[12] Our industry fixed effects are Fama-French 30 industries unless otherwise specified.[13] We cluster all standard errors at the firm level.

To facilitate comparing our results with the findings from prior studies on earlier time periods (i.e., Jung et al., 2018), we adopt a regression with two measures of *Earnings Surprise* as the variables of interest, along with their interaction, and restrict the event to only earnings announcement days. The first measure, *Missed Earnings*, is equal to one if earnings surprise is less than zero. The second measure is |*Earnings Surprise*|, the absolute value of *Earnings*

[12]As the number of fixed effects could cause some concern with using logistic regression, we replicated all our main findings using OLS regressions. The results are tabulated in Tables S8–S11 in Appendix S1.

[13]As our dependent variables in most regressions are sparse, including 30 industry dummies can lead to a lack of convergence in some cases. In such cases, we replace the Fama-French 30 industry fixed effects with Global Industry Classification Standard (GICS) sector fixed effects. All our results are robust to using GICS sector fixed effects throughout.

*Surprise*. In the interaction, the absolute value measure is demeaned, indicated as $|Earnings\ Surprise|_{\perp}$. The regression is shown in Equation (3):

$$\begin{aligned}\Phi^{-1}(Financial\ Tweets_{i,d}) = \alpha &+ \beta_1 Missed\ Earnings + \beta_2 |Earnings\ Surprise| \\ &+ \beta_3 Missed\ Earnings \times |Earnings\ Surprise|_{\perp} \\ &+ \beta_4 Financial\ Controls_{i,d} + Fixed\ Effects_{i,d}.\end{aligned} \tag{3}$$

In this specification, $\beta_2$ is of particular interest, as a positive coefficient would be consistent with our hypothesis. The coefficients $\beta_1$ and $\beta_3$ are also of interest. Negative coefficients would suggest that firms are less likely to post financial tweets when earnings news is significantly bad. In contrast, positive coefficients for $\beta_1$ and $\beta_3$ would support our hypothesis in the context of negative news earnings events. In this specification, we include only financial controls, and we use industry, year, and quarter fixed effects following Jung et al. (2018).

To bridge the gap between the abnormal return–based measures used in regression Equation (2) and the earnings surprise measures used in regression Equation (3), we mimic regression Equation (3) using abnormal return in place of earnings surprise. Specifically, we replace *Missed Earnings* with *Negative AR*, $|Earnings\ Surprise|$ with $|AR|$, and $Missed\ Earnings \times |Earnings\ Surprise|_{\perp}$ with $Negative\ AR \times |AR|$. *AR* is the continuous return used to calculate return quintiles for Equation (2).

## 5 | EMPIRICAL RESULTS

### 5.1 | Univariate results

Table 2 presents the summary statistics of our variables. The sample consists of 2,278,730 daily observations, where 63.9% of firm-days involve at least one tweet, and 3.2% (4.2%) of firm-days involve a financial tweet based on our ML (dictionary) measure. Untabulated results show that non-financial business content is present on 60.6% of firm-days, split between tweets about customer support (19.2%), conferences or tradeshows (40.9%), other marketing-related content (39.3%), and other business-related content (28.3%). Firms tweet content that is not related to business on 30.6% of firm-days. Overall, there is a large amount of overlap across the types of content posted each day, as firms post an average of 10.8 tweets per day. Verified Twitter accounts tend to be older and represent 37.2% of total observations. The number of followers and accounts followed are highly skewed, as the median (mean) observation has 5,839 (225,864) followers and is following 544 (2,731) accounts. Likewise, tweeting activity tends to be skewed, as the median (mean) observations have 1,897 (12,452) tweets in total. In addition, Panel B of Table 2 presents summary statistics for abnormal return quintiles. The mean abnormal return in Quintile 1 (5) is −0.026 (0.025), which suggests that these two quintiles consist of material negative (positive) news events.

We also examine the correlations between key dependent and independent variables in Panel C of Table 2. Financial tweets are positively correlated with earnings announcements and 10-K/10-Q filings. Although all presented correlations are statistically significant, the magnitude is generally smaller than 0.1. Lastly, Panel D of Table 2 presents summary statistics for the measures used to test Equation (3). Here we see that 27.6% of earnings announcements involved a negative earnings surprise, and that the average earnings surprise is positive. When we use abnormal return measures, we find that 48.8% of earnings announcements have a negative abnormal return on the announcement day.

Figure 1 presents the distribution of tweets and events within the week. The figure shows that all tweets and financial tweets are most prevalent when markets are open. For financial

**TABLE 2** Descriptive statistics.

**Panel A: Summary statistics, full sample**

| Variables | Mean | Median | SD | P10 | P90 |
|---|---|---|---|---|---|
| *Tweets* | 10.8 | 1 | 147 | 0 | 12.0 |
| *I*(*Tweets*) | 0.639 | 1.0 | 0.480 | 0 | 1 |
| *Financial Tweets, ML* | 0.0486 | 0 | 1.37 | 0 | 0 |
| *I*(*Financial Tweets, ML*) | 0.0319 | 0 | 0.176 | 0 | 0 |
| *Financial Tweets, Dict* | 0.439 | 0 | 1.97 | 0 | 0 |
| *I*(*Financial Tweets, Dict*) | 0.143 | 0 | 0.200 | 0 | 0 |
| *Earnings Ann* | 0.0158 | 0 | 0.125 | 0 | 0 |
| *Form 10-K, 10-Q* | 0.0154 | 0 | 0.123 | 0 | 0 |
| *Non-Earnings 8-K* | 0.114 | 0 | 0.318 | 0 | 1 |
| *AR* | −0.0005 | −0.0003 | 0.0529 | −0.019 | 0.0179 |
| *Verified* | 0.372 | 0 | 0.483 | 0 | 1 |
| *Followers* | 225,864 | 5,839 | 2,157,665 | 422 | 182,715 |
| *Friends* | 2,731 | 544 | 15,597 | 48.0 | 4,054 |
| *Recent Tweets* | 53.9 | 8.00 | 522 | 0 | 62.0 |
| *Total Tweets* | 12,452 | 1,897 | 74,359 | 91.0 | 15,989 |
| *Size* | 8.28 | 8.16 | 1.77 | 6.11 | 10.6 |
| *ROA* | 0.0113 | 0.0114 | 0.0364 | −0.0082 | 0.0362 |
| *MB* | 1.59 | 1.08 | 1.73 | 0.326 | 3.35 |
| *Debt* | 0.578 | 0.575 | 0.262 | 0.262 | 0.860 |
| *Growth* | 1.11 | 1.04 | 4.46 | 0.880 | 1.26 |
| *Advertising* | 0.0144 | 0.0000 | 0.0381 | 0.0000 | 0.0422 |
| *Return Volatility* | 0.0214 | 0.0178 | 0.0144 | 0.0104 | 0.0356 |
| *Firm Age* | 31.3 | 25.0 | 18.5 | 11.0 | 63.0 |
| Log(# *Analysts*)$_{\perp}$ | 0.0498 | 0.298 | 1.12 | −1.56 | 1.39 |
| *CEO Age* | 57.0 | 57.0 | 7.00 | 49.0 | 65.0 |

**Panel B: Summary statistics of each abnormal return quintile**

| Quintile | Mean | Median | SD | P10 | P90 |
|---|---|---|---|---|---|
| 1 (bottom) | −0.0262 | −0.0192 | 0.110 | −0.0467 | −0.0099 |
| 2 | −0.0074 | −0.0066 | 0.0040 | −0.0128 | −0.0032 |
| 3 | −0.0004 | −0.0003 | 0.0021 | −0.0031 | 0.0020 |
| 4 | 0.0065 | 0.0059 | 0.0032 | 0.0031 | 0.0110 |
| 5 (top) | 0.0248 | 0.0181 | 0.0282 | 0.0098 | 0.0440 |

**Panel C: Correlations between dependent and independent variables**

| | | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
|---|---|---|---|---|---|---|---|---|---|---|---|
| (1) | *I*(*Financial Tweets, ML*) | | | | | | | | | | |
| (2) | *I*(*Financial Tweets, Dict*) | 0.227 | | | | | | | | | |
| (3) | *Earnings Ann* | 0.055 | 0.046 | | | | | | | | |
| (4) | *Form 10-K, 10-Q* | 0.022 | 0.022 | 0.119 | | | | | | | |
| (5) | *Non-Earnings 8-K* | 0.012 | 0.029 | 0.156 | 0.085 | | | | | | |

**TABLE 2** (Continued)

**Panel C: Correlations between dependent and independent variables**

| | | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
|---|---|---|---|---|---|---|---|---|---|---|---|
| (6) | *Extreme Negative AR × Earnings Ann* | 0.039 | 0.032 | 0.554 | 0.087 | 0.061 | | | | | |
| (7) | *Non-extreme AR × Earnings Ann* | 0.017 | 0.014 | 0.595 | 0.030 | 0.140 | −0.005 | | | | |
| (8) | *Extreme Positive AR × Earnings Ann* | 0.039 | 0.034 | 0.573 | 0.090 | 0.066 | −0.005 | −0.006 | | | |
| (9) | *Extreme Negative AR × Form 10-K, 10-Q* | 0.016 | 0.015 | 0.093 | 0.526 | 0.057 | 0.176 | −0.005 | −0.005 | | |
| (10) | *Non-extreme AR × Form 10-K, 10-Q* | 0.007 | 0.007 | 0.026 | 0.657 | 0.037 | −0.006 | 0.054 | −0.006 | −0.005 | |
| (11) | *Extreme Positive AR × Form 10-K, 10-Q* | 0.016 | 0.016 | 0.099 | 0.530 | 0.057 | −0.005 | −0.005 | 0.181 | −0.004 | −0.006 |

**Panel D: Summary statistics of Table 5 measures**

| Variables | *N* | Mean | Median | SD | P10 | P90 |
|---|---|---|---|---|---|---|
| *Missed Earnings* | 18,722 | 0.276 | 0 | 0.447 | 0 | 1 |
| *Earnings Surprise* | 18,722 | 0.0004 | 0.0005 | 0.243 | −0.0018 | 0.0045 |
| $\lvert$*Earnings Surprise*$\rvert$ | 18,722 | 0.0092 | 0.0010 | 0.243 | 0.0001 | 0.0066 |
| $\lvert$*Earnings Surprise*$\rvert_{\perp}$ | 18,722 | −0.0062 | −0.0144 | 0.243 | −0.0153 | −0.0088 |
| *Negative AR* | 35,976 | 0.488 | 0 | 0.500 | 0 | 1 |
| $\lvert AR \rvert$ | 35,976 | 0.0311 | 0.0166 | 0.0459 | 0.0025 | 0.0752 |

*Note*: The sample consists of 2,278,730 observations. Panel A shows univariate statistics for the whole sample, while Panel B shows Pearson correlations between the dependent and independent variables. Correlation significance is omitted as all included correlations are significant at $p < 0.01$ with a Bonferroni adjustment. Panel C shows market model abnormal return statistics by quintile of market model abnormal return (calculated at the firm level). Panel D shows summary statistics for measures used only for testing regression Equation (3). Variable definitions for all variables included in subsequent regressions are included in Appendix 1. Methodology for the ML financial tweet measures is discussed in Section 3.2 and Appendix 2, and methodology for the dictionary financial tweets is discussed in Section 3.2. *I*() denotes an indicator variant of a measure, equal to one whenever the underlying measure is positive, and zero otherwise.

tweets there is a stronger concentration among NYSE opening hours, with a faster drop after the market closes and very little activity on weekends. For events, we see that earnings announcements are heavily concentrated within 1 h after the NYSE closes, as well as in weekday pre-market hours. 10-K/10-Q filings are more concentrated during trading hours than in pre-market hours, but they still exhibit a large spike after market close. Non-earnings 8-K filings show a pronounced spike after market close but also have a sizeable amount of disclosure before and during market hours. Figures S1 and S2 in Appendix S1 present complimentary visual evidence on the distribution of tweets on event days, arranging tweets by hour of release

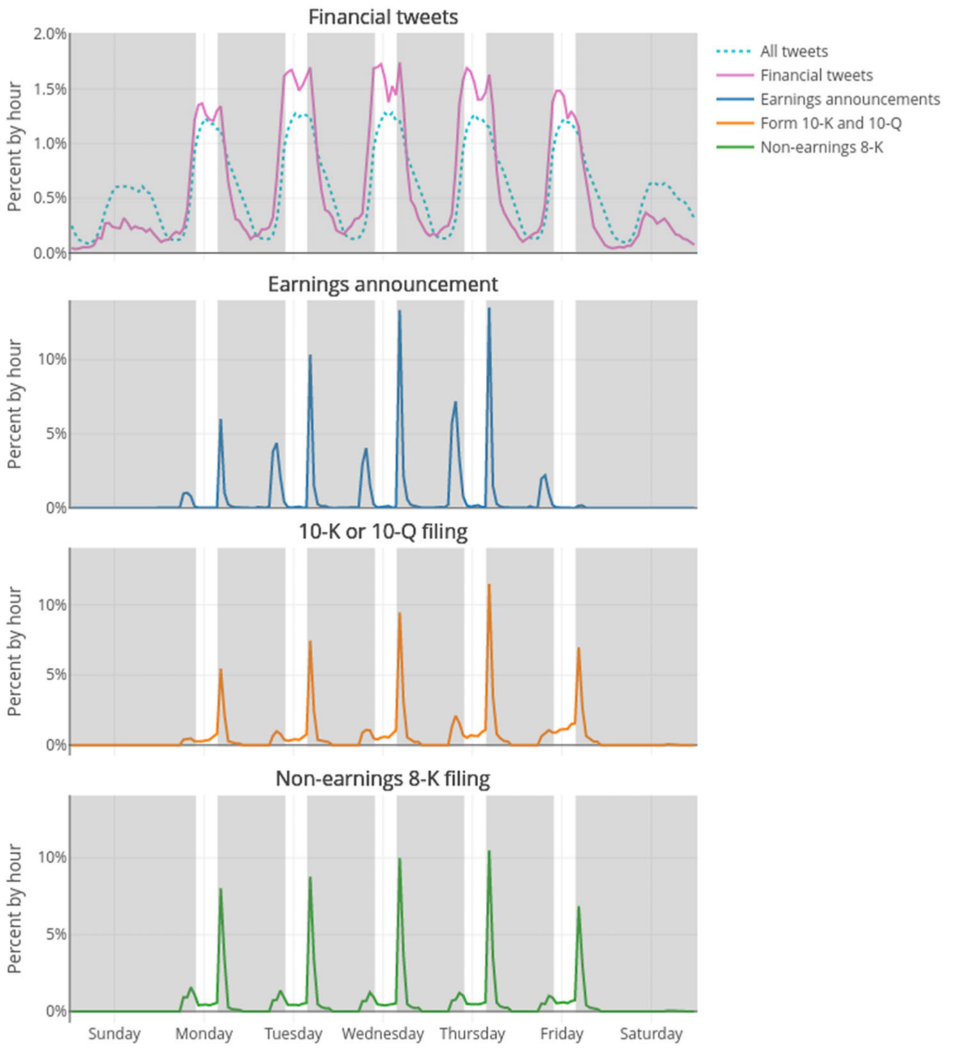


**FIGURE 1** Distribution of tweets and events by hour of the week. The top panel of this figure shows the distribution of firms' tweets, broken into all tweets (dotted blue) and financial tweets (solid pink, using our ML measure detailed in Section 3.2). The remaining three panels show the distribution of earnings announcements (blue), 10-K/10-Q filings (orange), and non-earnings announcement 8-K filings (green), respectively. The background is white during hours when the NYSE is open and gray when it is closed.

and time difference between the tweets and events, respectively. We see that financial tweets respond strongly to the presence of events.

## 5.2 | Tests of our hypothesis

Equation (1) examines the discretionary choice of tweeting financial information on event days. Table 3 presents the results. We find an increase in financial information dissemination on Twitter around earnings announcements and around 10-K/10-Q filings for both the ML and dictionary measures. However, the number of financial tweets does not increase around non-earnings 8-K filings when we use our ML measure. Using the dictionary measure, we find an

**TABLE 3** Association between financial tweets and events.

| ***I*(*Financial Tweets*) measure** | ***ML*** | | ***Dict*** | |
|---|---|---|---|---|
| **Variables** | **(1)** | ***Z*-stat** | **(2)** | ***Z*-stat** |
| *Earnings Ann* | 1.408*** | 18.77 | 0.835*** | 18.98 |
| *Form 10-K, 10-Q* | 0.562*** | 10.56 | 0.355*** | 11.22 |
| *Non-Earnings 8-K* | 0.006 | 0.16 | 0.149*** | 9.24 |
| *Verified* | 0.314** | 2.30 | 0.119* | 1.80 |
| Log(*Followers*) | 0.121** | 2.45 | −0.029 | −1.09 |
| Log(*Friends*) | −0.107*** | −2.68 | 0.002 | 0.12 |
| *Recent Tweets* | 0.000*** | 4.51 | — | |
| Log(*Total Tweets*) | 0.458*** | 7.28 | 0.559*** | 17.71 |
| *Size* | −0.040 | −0.84 | 0.055** | 2.17 |
| *ROA* | 3.158*** | 3.47 | 0.284 | 0.68 |
| *MB* | 0.065** | 2.44 | 0.032* | 1.67 |
| *Debt* | −0.142 | −0.57 | 0.146 | 1.36 |
| *Growth* | 0.006* | 1.91 | −0.001 | −0.54 |
| *Advertising* | −10.651*** | −3.64 | −4.908*** | −4.54 |
| *Return Volatility* | −0.975 | −0.30 | −2.088 | −1.41 |
| *Firm Age* | 0.000 | 0.02 | 0.001 | 0.75 |
| Log(*# Analysts*)$_{\perp}$ | −0.036 | −0.61 | 0.022 | 0.04 |
| *CEO Age* | 0.002 | 0.23 | −0.004 | 0.84 |
| Constant | −7.680*** | −12.24 | −5.031*** | −1.03 |
| Industry FE | Yes | | Yes | |
| Year FE | Yes | | Yes | |
| Month FE | Yes | | Yes | |
| Pseudo $R^2$ | 0.171 | | 0.147 | |
| Sample size | 2,278,730 | | 2,278,730 | |

*Note*: The regressions are run on the full sample. The dependent variable for all regressions is *I*(*Financial Tweets*), an indicator for whether a firm posted a financial tweet on a given trading day. Column 1 uses the ML-based financial tweets measure, while Column 2 uses the dictionary financial tweets measure. Variable definitions for all variables are included in Appendix 1. Methodology for the ML financial tweets measure is discussed in Section 3.2 and Appendix 2, while methodology for the dictionary approach is discussed in Section 3.2. All standard errors are clustered at the firm level. Industry fixed effects are based on Fama-French 30 industries. Variables dropped to ensure convergence are denoted by "—." *Z*-statistics are presented.
***, **, and * denote the significance levels for all coefficients at $p < 0.01$, $p < 0.05$, and $p < 0.10$, respectively.

increase in financial tweets around non-earnings 8-K filings. This is consistent with the dictionary approach capturing marketing and other business activities beyond the scope of financial information. As earnings announcements and 10-K/10-Q filings demonstrate a reaction using both measures, we focus all subsequent tests on these events.

The regressions based on Equation (2) are presented in Table 4. Columns 1 and 2 (3 and 4) present results for ML (dictionary) financial tweets measures. Each of our three variables of interest, *Extreme Negative AR*, *Extreme Positive AR*, and *Non-extreme AR*, is interacted with indicators for earnings announcements (Columns 1 and 3) or 10-K/10-Q filings (Columns 2 and 4). Our hypothesis predicts positive and significant coefficients for *Extreme Negative AR* and *Extreme Positive AR* that are higher than the coefficient on *Non-extreme AR*. In other words, we should observe a convex U-shaped relation between abnormal returns and the likelihood of posting a financial tweet. Consistent with our hypothesis, we see significant and positive coefficients on *Extreme Negative AR* and *Extreme Positive AR* (1.77 and 1.71, respectively) in

**TABLE 4** Financial tweets and materiality of events.

| ***I(Financial Tweets)* measure** | ***ML*** | | ***Dict*** | |
|---|---|---|---|---|
| ***Event*** | ***Earnings Ann*** | ***Form 10-K, 10-Q*** | ***Earnings Ann*** | ***Form 10-K, 10-Q*** |
| **Variables** | **(1)** | **(2)** | **(3)** | **(4)** |
| *Extreme Negative AR* × *Event* | 1.768*** | 1.048*** | 1.109*** | 0.682*** |
| | (21.95) | (14.67) | (21.26) | (14.29) |
| *Non-extreme AR* × *Event* | 0.899*** | 0.385*** | 0.573*** | 0.206*** |
| | (12.03) | (7.46) | (13.40) | (6.91) |
| *Extreme Positive AR* × *Event* | 1.709*** | 1.045*** | 1.119*** | 0.712*** |
| | (20.96) | (12.52) | (21.63) | (14.40) |
| *Verified* | 0.313** | 0.313** | 0.118* | 0.062 |
| | (2.29) | (2.30) | (1.78) | (0.94) |
| Log(*Followers*) | 0.120** | 0.120** | −0.029 | −0.025 |
| | (2.45) | (2.45) | (−1.10) | (−0.91) |
| Log(*Friends*) | −0.107*** | −0.107*** | 0.003 | 0.005 |
| | (−2.68) | (−2.69) | (0.13) | (0.26) |
| *Recent Tweets* | 0.000*** | 0.000*** | — | — |
| | (4.52) | (4.48) | | |
| Log(*Total Tweets*) | 0.458*** | 0.456*** | 0.452*** | 0.580*** |
| | (7.29) | (7.29) | (9.80) | (18.16) |
| *Size* | −0.040 | −0.040 | 0.057** | 0.051* |
| | (−0.84) | (−0.85) | (2.23) | (1.96) |
| *ROA* | 3.160*** | 3.130*** | 0.266 | 0.075 |
| | (3.47) | (3.45) | (0.64) | (0.18) |
| *MB* | 0.065** | 0.064** | 0.032* | 0.022 |
| | (2.44) | (2.43) | (1.66) | (1.20) |
| *Debt* | −0.142 | −0.140 | 0.148 | 0.025 |
| | (−0.56) | (−0.56) | (1.37) | (0.25) |
| *Growth* | 0.006* | 0.006* | −0.001 | −0.001 |
| | (1.91) | (1.95) | (−0.52) | (−0.98) |
| *Advertising* | −10.643*** | −10.589*** | −4.900*** | −3.253*** |
| | (−3.64) | (−3.64) | (−4.54) | (−3.88) |
| *Return Volatility* | −1.032 | −1.323 | −1.967 | −1.760 |
| | (−0.32) | (−0.41) | (−1.33) | (−1.17) |
| *Firm Age* | 0.000 | 0.000 | 0.001 | −0.001 |
| | (0.01) | (0.03) | (0.71) | (−0.31) |
| Log(*# Analysts*)$_{\perp}$ | −0.036 | −0.036 | 0.022 | 0.017 |
| | (−0.60) | (−0.61) | (0.85) | (0.64) |
| *CEO Age* | 0.002 | 0.001 | −0.004 | −0.005 |
| | (0.24) | (0.23) | (−1.03) | (−1.34) |
| Constant | −7.684*** | −7.607*** | −6.178*** | −6.051*** |
| | (−12.26) | (−12.19) | (−16.90) | (−19.53) |
| Industry FE | Yes | Yes | Yes | Yes† |
| Year FE | Yes | Yes | Yes | Yes |

**TABLE 4** (Continued)

| ***I(Financial Tweets)* measure** | ***ML*** | | ***Dict*** | |
|---|---|---|---|---|
| ***Event*** **Variables** | ***Earnings Ann*** **(1)** | ***Form 10-K, 10-Q*** **(2)** | ***Earnings Ann*** **(3)** | ***Form 10-K, 10-Q*** **(4)** |
| Month FE | Yes | Yes | Yes | Yes |
| $\chi^2$ test: *Negative* − *Non-extreme* | 168.67*** | 94.42*** | 162.93*** | 113.49*** |
| $\chi^2$ test: *Positive* − *Non-extreme* | 147.05*** | 72.35*** | 137.86*** | 120.94*** |
| $\chi^2$ test: *Negative* − *Positive* | 1.42 | 0.00 | 1.37 | 0.57 |
| Pseudo $R^2$ | 0.171 | 0.165 | 0.147 | 0.140 |
| Sample size | 2,278,730 | 2,278,730 | 2,278,730 | 2,278,730 |

*Note*: The regressions are run on the full sample. The dependent variable for all regressions is *I*(*Financial Tweets*), an indicator for whether a firm posted a financial tweet on a given trading day. Columns 1 and 2 use the ML-based financial tweets measure, while Columns 3 and 4 use the dictionary financial tweets measure. *Event* is an indicator for earnings announcement days in Columns 1 and 3 and for 10-K/10-Q filing days in Columns 2 and 4. Variable definitions for all variables are included in Appendix 1. Methodology for the ML financial tweets measure is discussed in Section 3.2 and Appendix 2, while methodology for the dictionary approach is discussed in Section 3.2. All standard errors are clustered at the firm level. Industry fixed effects are based on Fama-French 30 industries, unless denoted with a dagger (†), in which case industry fixed effects are based on GICS sector (due to lack of convergence from the large number of measures generated for Fama-French 30 industries). Chi-squared tests denote tests of significant difference between the independent variables—for example, *Negative* − *Non-extreme* denotes a two-tailed test for the difference between *Extreme Negative AR* × *Event* and *Non-extreme AR* × *Event*. Variables dropped to ensure convergence are denoted by "—." *Z*-statistics are presented in parentheses.

***, **, and * denote the significance levels for all coefficients at $p < 0.01$, $p < 0.05$, and $p < 0.10$, respectively.

Column 1. Furthermore, these coefficients are larger than that of *Non-extreme AR* (0.90), and thus there does appear to be a convex U-shaped relation between the materiality of earnings announcements and financial tweet posting. Column 2 of Table 4 shows the result of testing our hypothesis around 10-K/10-Q filings, again demonstrating the expected convex U-shaped relation. Using chi-squared tests, we confirm that the response to extreme negative and extreme positive events is larger than the response to lower-materiality events for both earnings announcements and 10-K/10-Q filings. In addition, we test whether the extreme negative and positive responses are different from each other and from the lower-materiality responses. We find no statistical difference between extremes for both events, which suggests no deviation from a symmetric relation.[14]

In Columns 3 and 4, we present the results with the dictionary financial tweets as the dependent variable. Our results are robust. We find positive and significant coefficients for *Extreme Negative AR* and *Extreme Positive AR* for both events. As with the ML measure, we find no significant differences between *Extreme Negative AR* and *Extreme Positive AR* for either event. Furthermore, we find statistically significant differences between the extreme abnormal return measures and the non-extreme measures for both earnings announcements and 10-K/10-Q filings.

Figure 2 illustrates a more granular version of our findings on our hypothesis. The figure follows regression Equation (2), but excludes abnomal return indicators, instead presenting the average residuals of the model in 1% increments of abnormal return percentiles around events and for each financial tweet measure. Figure 2A presents the results using the ML financial tweets measure. Even at such a granular level, there is a general convex U-shape to the model residuals, with the highest residuals occurring at the lowest and highest percentiles, and the

[14]In Table S12 in Appendix S1, we rerun our analyses including interactions for earnings announcements, 10-K/10-Q filings, and non-earnings 8-K filings in the same regression to replicate Table 4. All coefficients on earnings announcements and 10-K/10-Q filings remain significant at $p < 0.01$, and we observe the same U-shaped relation for both event types.

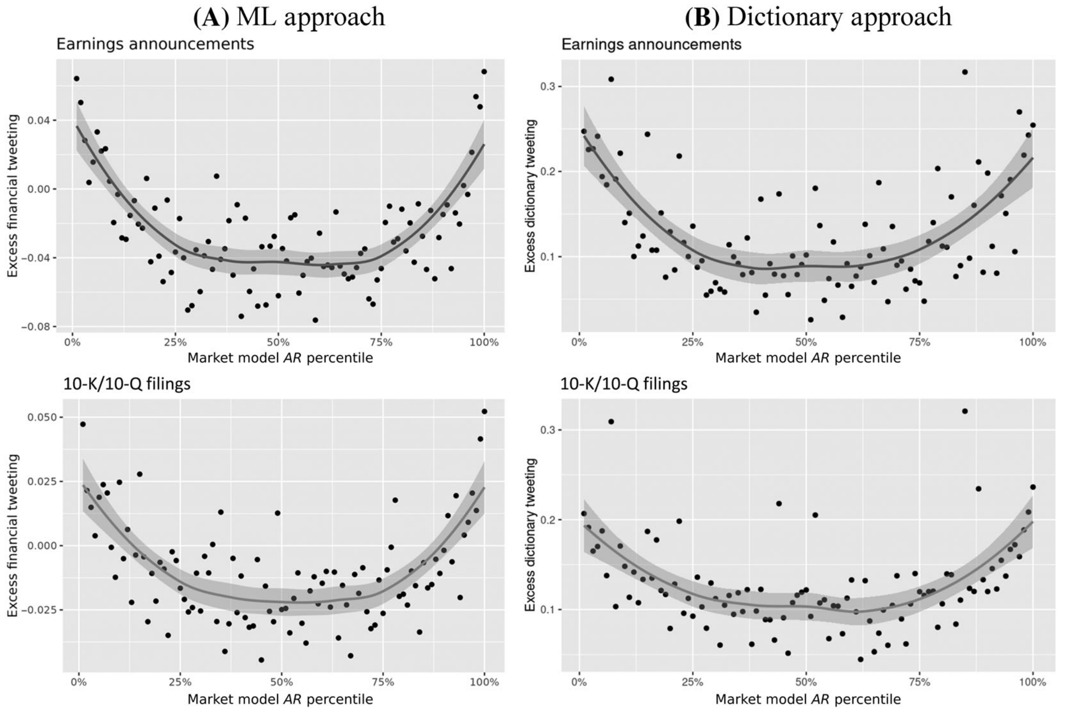


**FIGURE 2** U-shaped relation between financial tweets and materiality of earnings announcements and 10-K/10-Q filings: (A) ML approach and (B) dictionary approach. The dots are regression residuals following the regression structure reported in Table 4, excluding extreme and non-extreme abnormal return indicators, and grouping by abnormal return percentiles with a 1% increment. The trend line on each plot is computed using locally estimated scatterplot smoothing (LOESS) with a 95% confidence interval around the trend line. Financial tweets are based on our ML measure detailed in Section 3.2.

lowest residuals occurring in the middle of the abnormal return distribution. For 10-K/10-Q filings, we again document a pronounced convex U-shaped relation between materiality and financial tweeting. Figure 2B presents the results using the dictionary financial tweet measure. We again see a pronounced U-shape for both earnings announcements and 10-K/10-Q filings.

## 6 | SENSITIVITY ANALYSES

The documented U-shaped relation differs from the findings in Jung et al. (2018), which show that managers are less likely to disseminate bad news in the setting of earnings announcements. We show that for both extreme negative and positive events (that lead to an abnormal return in the bottom or top 20% of the firm's abnormal return distribution), firms are more likely to disseminate financial information on Twitter than they are around less material events. The results are robust to using both a more precise ML-based measure of financial tweets and a dictionary measure built from the word list in Jung et al. (2018). To further understand the cause of the difference, that is, whether the strategic behavior of managers on social media has evolved over time, we conduct three additional analyses: (1) using a regression specification that more closely follows Jung et al. (2018), (2) examining the intraday pattern of financial tweets around extreme negative and positive events, and (3) partitioning the sample by the year firms joined Twitter.

## 6.1 | Regression with a specification similar to that of Jung et al. (2018)

Equation (3) describes the model specification of Jung et al. (2018), which uses earnings surprise to label events.[15] We tabulate the results using two different sample approaches. The first approach follows Jung et al. (2018) and restricts the test to earnings announcement days (EA sample). The second approach uses our full daily sample of tweets (full sample), using non-event days to control for baseline tweeting behavior outside of earnings announcement days.

In Table 5, Columns 1 and 2 present the results using the EA sample. We see that the coefficients on both *Missed Earnings* and |*Earnings Surprise*| are insignificant. Consistent with our earlier results, we find a positive and significant coefficient on the interaction term using the ML measure, but we do not find any result using the dictionary measure.[16] This result contrasts with the significant negative interaction term found in Jung et al. (2018), which has a different sample size and time period.

We present results using the full sample in Columns 3 and 4. Using our ML measure of financial tweets (Column 3), we find significantly positive coefficients for *Missed Earnings*, |*Earnings Surprise*|, and the interaction term. Using the dictionary measure of financial tweets (Column 4), we find positive and significant results for *Missed Earnings* and the interaction term. In contrast to our ML measure results, we find insignificant coefficients for |*Earnings Surprise*|. Thus, only with our ML measure do we continue to observe results consistent with a U-shaped pattern of disclosure. The difference in the coefficients across the two measures highlights the importance of using a more precise measure for the tests. The positive sign on *Missed Earnings* in both specifications suggests that firms are more likely to post financial tweets when earnings news is negative. This implies that the tweeting pattern in the later period (our sample) is different from the pattern documented in the earlier period in Jung et al. (2018).

Furthermore, we restructure regression Equation (3) to use measures based on abnormal returns. Specifically, we replace *Missed Earnings* with *Negative AR* and |*Earnings Surprise*| with $|AR|$. The results of these tests are presented in Table 5, Columns 5–8. Columns 5 and 6 show a positive and significant coefficient on $|AR|$, suggesting a U-shaped relation for the EA sample. When we expand to the full sample in Columns 7 and 8, we continue to find a positive and significant coefficient on $|AR|$ for both financial tweet measures. Additionally, the coefficients on the interaction term are significantly positive when we use our ML measure, again suggesting that firms are more likely to post financial tweets when earnings news is negative. In contrast, using the dictionary measure leads to a negative coefficient on the interaction. This negative coefficient is driven by the investor-focused terms in the dictionary, as including only earnings-related terms leads to a positive and significant interaction coefficient (see Table S4 in Appendix S1).

## 6.2 | Intraday analysis

In this subsection, we examine the intraday timing of tweets as an exploratory analysis. Financial tweets are almost entirely concurrent with or follow earnings announcements. For 10-K/10-Q filings, there is a build-up of financial tweets in the 12 h beforehand, along with a peak in the 3 h around the filing release. Given these relations, we examine how hourly financial tweeting changes using an hourly sample spanning the 48 h around our events.

[15]All controls follow Jung et al. (2018). We do not have three control variables included in tab. 3, panel B, of Jung et al. (2018) due to lack of data: the number of press releases firms issued, the number of news articles covering the firm, and an indicator for whether a firm is headquartered in Silicon Valley. Our results are robust to including the count of 8-Ks released during a quarter as a proxy for the number of press releases issued by the firm.

[16]The lack of any result for the dictionary measure is driven by the inclusion of investor-focused terms. If we only include earnings-related terms, we find a positive and significant coefficient on the interaction term (see Table S4, Column 1, in Appendix S1). If we only include investor-focused terms, we find no significant results (untabulated).

**TABLE 5** Regression with a specification similar to that of Jung et al. (2018).

| Model | EA sample | | Full sample | | EA sample, *AR* | | Full sample, *AR* | |
|---|---|---|---|---|---|---|---|---|
| ***I(Financial Tweets)* measure** | ***ML*** | ***Dict*** | ***ML*** | ***Dict*** | ***ML*** | ***Dict*** | ***ML*** | ***Dict*** |
| **Variables** | **(1)** | **(2)** | **(3)** | **(4)** | **(5)** | **(6)** | **(7)** | **(8)** |
| *Missed Earnings* | −0.084 | −0.004 | 1.345*** | 0.797*** | | | | |
| | (−1.16) | (−0.09) | (11.64) | (11.94) | | | | |
| $\lvert$*Earnings Surprise*$\rvert$ | −0.022 | −0.168 | 0.079* | −0.006 | | | | |
| | (−0.26) | (−0.42) | (1.78) | (−0.11) | | | | |
| *Missed Earnings* × $\lvert$*Earnings Surprise*$\rvert_{\perp}$ | 0.794*** | 0.394 | 0.366** | 0.153* | | | | |
| | (4.30) | (1.06) | (2.35) | (1.92) | | | | |
| *Negative AR* | | | | | 0.000 | −0.026 | 0.001 | 0.004 |
| | | | | | (0.00) | (−0.80) | (0.08) | (0.66) |
| $\lvert AR\rvert$ | | | | | 5.558*** | 3.485*** | 5.008*** | 1.744*** |
| | | | | | (3.93) | (5.57) | (6.65) | (4.09) |
| *Negative AR* × $\lvert AR\rvert$ | | | | | 0.829 | 0.319 | 1.589*** | −0.584** |
| | | | | | (0.64) | (0.56) | (2.66) | (−2.09) |
| *Size* | 0.386*** | 0.282*** | 0.220*** | 0.204*** | 0.449*** | — | 0.227*** | — |
| | (4.41) | (4.77) | (5.22) | (8.17) | (9.91) | | (5.41) | |
| *ROA* | 2.866** | 1.184 | 4.367*** | 1.074* | 2.971*** | 3.362*** | 4.535*** | 2.177*** |
| | (2.30) | (1.17) | (3.77) | (1.95) | (2.99) | (4.70) | (4.08) | (4.19) |
| *MB* | 0.087 | 0.092** | 0.150*** | 0.086*** | 0.125*** | 0.037 | 0.152*** | 0.046** |
| | (1.59) | (2.20) | (5.29) | (4.20) | (4.03) | (1.44) | (5.39) | (2.21) |
| *Debt* | 0.592 | 0.539** | 0.114 | 0.393*** | 0.233 | 0.908*** | 0.100 | 0.661*** |
| | (1.27) | (1.98) | (0.38) | (2.82) | (0.88) | (4.60) | (0.34) | (4.89) |
| *Growth* | −0.003 | 0.004 | 0.003 | −0.002 | 0.001 | 0.002 | 0.003 | −0.002*** |
| | (−1.36) | (1.24) | (0.89) | (−1.62) | (0.28) | (0.61) | (0.85) | (−2.66) |
| *Advertising* | 0.760 | −0.718 | −3.319* | −1.113 | −0.571 | −0.500 | −3.438** | −1.150 |
| | (0.39) | (−0.44) | (−1.93) | (−1.32) | (−0.39) | (−0.50) | (−1.99) | (−1.56) |

**TABLE 5** (Continued)

| **Model** | **EA sample** | | **Full sample** | | **EA sample, *AR*** | | **Full sample, *AR*** | |
|---|---|---|---|---|---|---|---|---|
| ***I*(*Financial Tweets*) measure** | ***ML*** | ***Dict*** | ***ML*** | ***Dict*** | ***ML*** | ***Dict*** | ***ML*** | ***Dict*** |
| **Variables** | **(1)** | **(2)** | **(3)** | **(4)** | **(5)** | **(6)** | **(7)** | **(8)** |
| *Firm Age* | 0.013** | 0.011*** | 0.002 | 0.004 | 0.008** | 0.015*** | 0.002 | 0.009*** |
| | (2.52) | (2.88) | (0.39) | (1.51) | (2.18) | (5.55) | (0.45) | (4.04) |
| Log(# *Analysts*) | 0.176 | 0.215* | 0.015 | 0.076** | 0.028 | 0.171*** | 0.018 | 0.154*** |
| | (0.99) | (1.80) | (0.23) | (2.47) | (0.52) | (4.39) | (0.27) | (4.72) |
| *CEO Age* | 0.003 | −0.004 | −0.008 | −0.009** | −0.001 | −0.005 | −0.008 | −0.008** |
| | (0.25) | (−0.53) | (−1.07) | (−2.10) | (−0.18) | (−0.94) | (−1.00) | (−2.02) |
| Constant | −6.398*** | −4.715*** | −6.306*** | −4.905*** | −7.165*** | −2.391*** | −6.459*** | −3.078*** |
| | (−6.84) | (−5.95) | (−10.12) | (−11.64) | (−11.08) | (−6.09) | (−10.41) | (−7.73) |
| Industry FE | Yes | Yes | Yes | Yes | Yes† | Yes† | Yes | Yes |
| Year FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Quarter FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Pseudo $R^2$ | 0.160 | 0.112 | 0.0978 | 0.0674 | 0.117 | 0.0479 | 0.0984 | 0.0560 |
| Sample size | 18,624 | 18,687 | 2,278,730 | 2,278,730 | 35,976 | 35,976 | 2,278,730 | 2,278,730 |

*Note*: The regressions in Columns 1, 2, 5, and 6 are run on a sample restricted to only earnings announcement days. Columns 1 and 2 are additionally restricted to only days with *Earnings Surprise* available in I/B/E/S. The dependent variable for all regressions is *I*(*Financial Tweets*), an indicator for whether a firm posted a financial tweet on a given trading day. Odd-numbered columns use the ML-based financial tweets measure, while even-numbered columns use the dictionary financial tweets measure. Variable definitions for all variables are included in Appendix 1. Methodology for the ML financial tweets measure is discussed in Section 3.2 and Appendix 2, while methodology for the dictionary approach is discussed in Section 3.2. All standard errors are clustered at the firm level. Industry fixed effects are based on Fama-French 30 industries, unless denoted with a dagger (†), in which case industry fixed effects are based on GICS sector (due to lack of convergence from the large number of measures generated for Fama-French 30 industries). Variables dropped to ensure convergence are denoted by "—." *Z*-statistics are presented in parentheses.
***, **, and * denote the significance levels for all coefficients at $p < 0.01$, $p < 0.05$, and $p < 0.10$, respectively.

For our intraday analysis, we examine 48-h periods around events, [−24, +24] h, using a firm-hour sample. Due to time-stamp data availability in I/B/E/S, our intraday tests have significantly fewer events than our other tests. We follow Equation (2), including the firm, year, and month fixed effects, but we use a more restrictive definition for the dependent variable and the event indicators. Our dependent variable is equal to one only when there is a financial tweet by the firm in a given hour. We present our results using event indicators equal to one for only the 3 h before (−3, 0) or the 3 h after (0, +3) the event occurred, and our results are robust to additional windows using 2 or 1 h before and after the events (untabulated). We include the interactions of these before- and after-event indicators with our materiality measures (*Extreme Negative AR*, *Non-extreme AR*, and *Extreme Positive AR*). We also add an additional fixed effect to this model: the hour at the NYSE. As the number of financial tweets varies by hour within the day, this hour-at-NYSE fixed effect controls for any natural variation in tweets due to the time of day. We cluster all standard errors at the firm level.

We first examine if there is an increased likelihood of financial tweet posting in the 3 h before and after our events. Table 6 presents the results. Using our ML measure of financial tweeting around earnings events, we find that financial tweeting is only concentrated in the 3 h after announcement, but we do not observe a U-shape in this shorter window. For 10-K/10-Q filings, we observe an elevated level of tweeting for extreme negative events and a U-shape in the 3 h after the event.

In an additional test presented in Figure S3 in Appendix S1, we expand the intraday test to include indicators for additional windows from 6 h before through 18 h after the event. Most financial tweets are posted within 6 h of the events. To gain further insight into how timing affects firms' use of Twitter, we partition this sample into events that occur before or after 4:00 p.m. We find that usage of Twitter is markedly different across these samples. The U-shaped relation is very visible in the 6 h following an earnings announcement posted after market close, whereas earlier announcements see a fairly even provision of financial tweets regardless of the materiality of the earnings announcement. For 10-K/10-Q filings, we see the U-shape take hold in both samples. These results suggest that there is interesting within-day variation in firm Twitter usage around these events.

### 6.3 | Partition analysis

Separately, we examine whether the firms that joined Twitter are different in earlier versus later years. We partition the 1,360 firms in our main sample into two groups: 901 early adopters who had already opened their accounts by the start of our sample period (2012) and 459 late adopters who opened their accounts during the sample period. We retest our main results across these two groups. Untabulated results replicating Table 4 are robust across both groups. However, we do find that the magnitude of the effects we document are larger for late adopters. For instance, the coefficient on negative (positive) news in the early-adopter group is 1.66 (1.61), while the same coefficient in the late-adopter group is 2.39 (2.31). This suggests that late adopters may use Twitter more strategically than early adopters.

## 7 | DESCRIPTIVE AND ADDITIONAL ANALYSES

### 7.1 | Trend in firm Twitter usage

The results documented in the previous section suggest that the relationship between financial tweeting and earnings announcements is U-shaped and skewed such that firms are more likely to tweet around extreme earnings news when the news is negative. The inconsistency between

**TABLE 6** Financial tweets before and after events: Intraday analysis.

| *I*(*Financial Tweets*) measure | Hourly, ML | | Hourly, Dict | |
|---|---|---|---|---|
| *Event* | *Earnings Ann* | *Form 10-K, 10-Q* | *Earnings Ann* | *Form 10-K, 10-Q* |
| **Variables** | **(1)** | **(2)** | **(3)** | **(4)** |
| *Extreme Negative AR* × [−3, 0) *Hours from Event* | 0.092 | 0.844*** | 0.239*** | 0.613*** |
| | (0.59) | (6.81) | (3.08) | (9.37) |
| *Non-extreme AR* × [−3, 0) *Hours from Event* | 0.046 | 0.455*** | 0.158* | 0.302*** |
| | (0.32) | (5.01) | (1.83) | (6.31) |
| *Extreme Positive AR* × [−3, 0) *Hours from Event* | 0.182 | 0.583*** | 0.166* | 0.514*** |
| | (1.14) | (4.90) | (1.78) | (8.15) |
| *Extreme Negative AR* × [0, +3) *Hours from Event* | 2.448*** | 1.218*** | 1.916*** | 0.963*** |
| | (22.06) | (10.38) | (24.59) | (14.72) |
| *Non-extreme AR* × [0, +3) *Hours from Event* | 2.475*** | 0.957*** | 2.025*** | 0.616*** |
| | (20.96) | (10.03) | (25.62) | (11.39) |
| *Extreme Positive AR* × [0, +3) *Hours from Event* | 2.472*** | 1.172*** | 1.994*** | 0.846*** |
| | (21.76) | (9.92) | (25.46) | (12.16) |
| *Verified* | 0.176 | 0.337** | 0.097 | 0.118 |
| | (1.08) | (2.38) | (0.90) | (1.35) |
| Log(*Followers*) | 0.001 | 0.077 | −0.060 | −0.049 |
| | (0.02) | (1.44) | (−1.51) | (−1.48) |
| Log(*Friends*) | −0.014 | −0.060 | −0.011 | 0.023 |
| | (−0.28) | (−1.45) | (−0.41) | (0.98) |
| *Recent Tweets* | 0.000 | 0.000** | 0.001*** | 0.000** |
| | (1.33) | (2.48) | (5.49) | (2.11) |
| Log(*Total Tweets*) | 0.205*** | 0.255*** | 0.366*** | 0.439*** |
| | (3.71) | (3.39) | (8.46) | (9.40) |
| *Size* | 0.327*** | 0.073 | 0.193*** | 0.140*** |
| | (5.12) | (1.40) | (4.26) | (4.34) |
| *ROA* | 0.887 | 3.458*** | 0.167 | 0.707 |
| | (0.77) | (3.20) | (0.23) | (1.06) |
| *MB* | 0.086** | 0.063** | 0.041 | 0.039* |
| | (2.44) | (2.37) | (1.50) | (1.90) |
| *Debt* | 0.067 | −0.012 | 0.297 | −0.024 |
| | (0.19) | (−0.05) | (1.49) | (−0.17) |
| *Growth* | −0.002 | −0.007*** | −0.000 | −0.003 |
| | (−1.14) | (−3.69) | (−0.06) | (−0.37) |
| *Advertising* | −1.447 | −6.949*** | −2.114 | −5.819*** |
| | (−0.72) | (−2.62) | (−1.36) | (−3.01) |
| *Return Volatility* | −5.193 | −1.384 | −0.007 | 2.152 |
| | (−1.34) | (−0.45) | (−0.00) | (1.22) |
| *Firm Age* | −0.002 | −0.001 | 0.002 | 0.003 |
| | (−0.52) | (−0.17) | (0.78) | (1.04) |
| Log(# *Analysts*)$_{\perp}$ | 0.123 | −0.032 | 0.077 | 0.001 |
| | (0.95) | (−0.55) | (0.84) | (0.04) |

(Continues)

**TABLE 6** (Continued)

| *I(Financial Tweets)* **measure** | *Hourly, ML* | | *Hourly, Dict* | |
|---|---|---|---|---|
| *Event* | *Earnings Ann* | *Form 10-K, 10-Q* | *Earnings Ann* | *Form 10-K, 10-Q* |
| **Variables** | **(1)** | **(2)** | **(3)** | **(4)** |
| *CEO Age* | −0.004 | −0.005 | −0.004 | −0.008* |
| | (−0.42) | (−0.66) | (−0.73) | (−1.70) |
| Constant | −12.266*** | −10.775*** | −9.720*** | −9.832*** |
| | (−12.14) | (−14.49) | (−16.59) | (−17.61) |
| Industry FE | Yes | Yes | Yes | Yes |
| Year FE | Yes | Yes | Yes | Yes |
| Month FE | Yes | Yes | Yes | Yes |
| Hour-at-NYSE FE | Yes | Yes | Yes | Yes |
| Window: [−3, 0) h | | | | |
| $\chi^2$ test: *Negative − Non-extreme* | 0.97 | 13.27*** | 0.90 | 31.83*** |
| $\chi^2$ test: *Positive − Non-extreme* | 0.04 | 1.40 | 0.01 | 14.92*** |
| $\chi^2$ test: *Negative − Positive* | 0.44 | 4.84** | 0.55 | 2.48 |
| Window: [0, +3) h | | | | |
| $\chi^2$ test: *Negative − Non-extreme* | 0.04 | 8.76*** | 5.26 | 39.81*** |
| $\chi^2$ test: *Positive − Non-extreme* | 1.17 | 6.15** | 0.36 | 18.21*** |
| $\chi^2$ test: *Negative − Positive* | 0.49 | 0.21 | 2.28 | 3.92** |
| Pseudo $R^2$ | 0.238 | 0.133 | 0.248 | 0.333 |
| Sample size | 954,432 | 1,689,888 | 954,432 | 1,689,888 |

*Note*: The regressions are run on an hourly sample restricted to 48-h intervals around earnings announcements (Columns 1 and 3) or 10-K/10-Q filings (Columns 2 and 4). Columns 1 and 3 require the earnings announcement to be included in I/B/E/S (for the time of the announcement). The dependent variable for all regressions is *I(Financial Tweets)*, an indicator for whether a firm posted a financial tweet on a given trading day. Columns 1 and 2 use the ML-based financial tweets measure, while Columns 3 and 4 use the dictionary financial tweets measure. The variable *[−3, 0) Hours from Event* is an indicator for the 3-h period leading up to the clock hour of the event; the variable *[0, +3) Hours from Event* is an indicator for the 3-h period starting from the calendar hour of the event. *Event* indicates earnings announcement days in Columns 1 and 3 and 10-K/10-Q filing days in Columns 2 and 4. Variable definitions for all variables are included in Appendix 1. Methodology for the ML financial tweets measure is discussed in Section 3.2 and Appendix 2, while methodology for the dictionary approach is discussed in Section 3.2. All standard errors are clustered at the firm level. Industry fixed effects are based on Fama-French 30 industries. Chi-squared tests denote tests of significant difference between the independent variables—for example, *Negative − Non-extreme* denotes a two-tailed test for the difference between *Extreme Negative AR × Event* and *Non-extreme AR × Event*. *Z*-statistics are presented in parentheses.
***, **, and * denote the significance levels for all coefficients at $p < 0.01$, $p < 0.05$, and $p < 0.10$, respectively.

these results and the findings in prior studies could be due to the difference in sample and a change in the strategic behaviors of managers over time. We thus further examine whether our results are consistent across different time periods or show a trend over time.

When we run the regression specified in Equation (1) separately by year, we find that financial tweeting is significantly associated with earnings announcements and 10-K/10-Q filings in each year (as summarized in Figure S4 in Appendix S1). Figure 3 presents the coefficients using the specification in Equation (2), except that the regressions are run for each year. The figure shows that firms react to both extreme good and bad news in earlier years of our sample period, but the U-shaped pattern is more pronounced in later years. Chi-squared tests confirm that the U-shape is statistically significant for each year in the sample for earnings announcements and in each year except 2012 and 2015 in the sample for 10-K/10-Q filings. This finding suggests that firms use Twitter more frequently to convey their financial information over time, although the U-shaped pattern persists throughout the entire sample period.

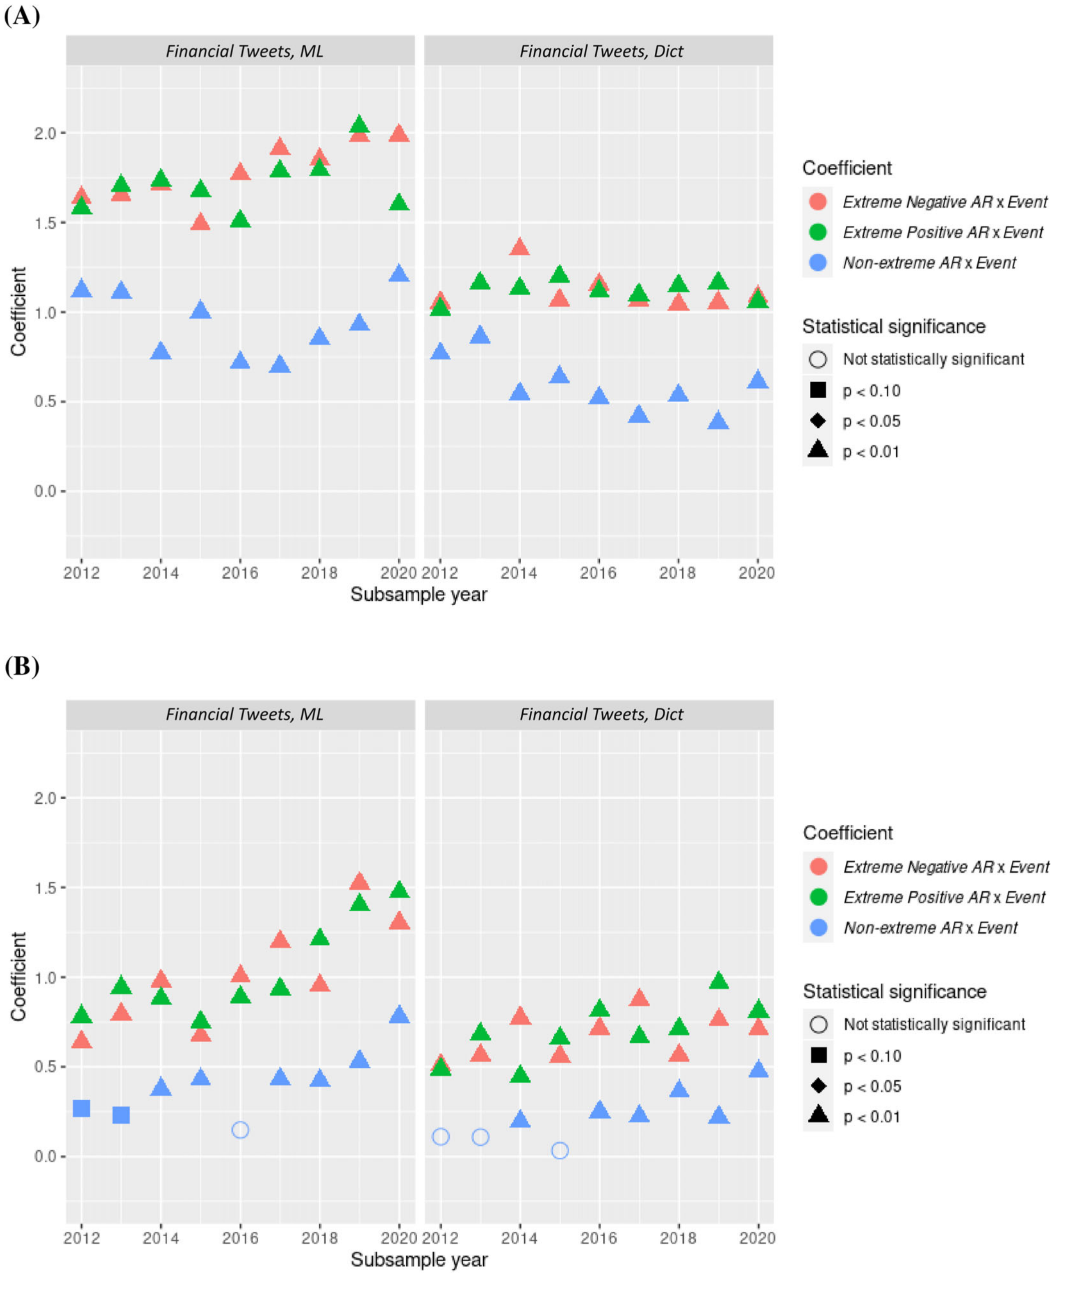


**FIGURE 3** Trend in firm Twitter usage by year: (A) earnings announcements and (B) 10-K/10-Q filings. This figure presents the coefficients and significance of the independent variables from Table 4, Columns 1 and 3, when running the regression for each year from 2012 to 2020.

## 7.2 | Firms' use of hyperlinks

The number of characters allowed in a tweet is limited, but embedded hyperlinks can bring investors' attention to the message tweeted (Nekrasov et al., 2022). If managers tweet extreme good and extreme bad news symmetrically, we expect that they are also likely to use hyperlinks in both good and bad news settings. We thus examine how firms' use of hyperlinks in financial tweets on Twitter varies with accounting events. On average, firms include a hyperlink in a financial tweet on 82.0% of all days with financial tweets (2.62% of all days, 80.1% of all financial tweets). This number suggests that over time, firms have become more likely to include hyperlinks in their tweets. Blankespoor et al. (2014) report that 75.4% of the tweets in their sample contain hyperlinks.

Replacing the dependent variable with an indicator of financial tweets containing hyperlinks in Equation (2) and restricting our sample to only days with financial tweets, we find results consistent with the expected U-shaped relation for both event types. The analyses, presented in Table S6 in Appendix S1, show significantly positive coefficients on both *Extreme Negative AR* and *Extreme Positive AR*. The results are robust using both the ML and dictionary measures of financial tweets and both event types, and the coefficients on *Extreme Negative AR* and *Extreme Positive AR* are significantly higher than the coefficient on *Non-extreme AR* in all but one specification each. The evidence that firms are more likely to disseminate financial information via hyperlinks within financial tweets around both material good and bad events further supports the existence of the strategic behavior of firms on social media.

## 7.3 | Determinants of using Twitter

We explore the factors driving the decision to create and use a firm Twitter account, as it is potentially important to control for these factors in our regressions. Our sample for this analysis comprises 1,321 firms with an active Twitter account and 514 firms that have no account or have never tweeted by the end of our sample period.[17] We run a logistic model across an annual version of our data set using our firm financial variables as controls along with institutional ownership (*Inst Ownership*), and we include year, month, and industry fixed effects. The regression results presented in Table S7 in Appendix S1, suggest that larger firms and firms with excess analyst following are more likely to have Twitter accounts and to have tweeted (*Tweeted*), as are those with more expected growth opportunities. Examining the trend in year fixed effects (using chi-squared tests, untabulated), we find a statistically significant increase in the number of firms that have joined Twitter each year from 2013 through 2020. This implies that the likelihood of a firm using Twitter consistently grew throughout our sample period. Examining industry fixed effects, we find the highest positive effects for firms in Printing and Publishing; Restaurants, Hotels, and Motels; Retail; and Personal and Business Services.

## 7.4 | Robustness checks

### 7.4.1 | Investor relations Twitter accounts

Some firms operate an investor relations (IR) Twitter account separate from their main account. In July 2017, we checked to see if each firm that had been in the S&P 1500 from 2012 to 2016 had a separate IR Twitter account. Among all firms, we found only 11 IR accounts. Furthermore, 862 of the firms we were tracking linked to their main Twitter account from their IR website, including eight of the firms with IR accounts. Only one firm linked to an IR Twitter account from its IR website. Thus, firms' main Twitter accounts appear to be the most important accounts for IR. Furthermore, our results for Tables 4 and 5 are robust across the

[17]The number of firms in each group in the determinants test is less than the number of firms in the sample described in Section 3.1 because we include institutional ownership in this test as a control. We drop 39 firms with active accounts, 44 firms without an account, and 12 firms that have accounts without tweets.

following subsamples: (1) removing the 11 firms with an IR Twitter account, (2) restricting the sample to the 862 firms that link from their IR website to their main Twitter account, and (3) restricting the sample to the 581 firms that do not link from their IR website to their main Twitter account.

### 7.4.2 | Econometric concerns

As some of the indicators we used in logistic regressions are interactions between multiple variables, we test the marginal effects of the interactions as in Norton et al. (2004). Our primary results in Tables 4 and 5 have consistent and significant marginal effects at the same significance level as our presented coefficients for every measure and interaction.

Second, as we use logistic regressions including a multitude of fixed effects for our main analyses, we replicate our analyses using OLS regression. A full replication of Tables 3–6 is available in Tables S8–S11 in Appendix S1. Focusing on our main results in Table 4, we find the same U-shaped pattern for both events. Replicating Table 5, we continue to find positive and significant coefficients for every measure that was positive and significant in the main tests, except the main effect of |*Earnings Surprise*| in Column 3.

### 7.4.3 | Tweet counts

As the dependent variables in our tests are derived from tweet counts, we examine the count of financial tweets in place of the incidence of such tweets. This measures the increase in financial tweeting, as opposed to the increased probability of financial tweeting, but we caution that financial tweet counts frequently equal one (approximately 77% of the time for both measures). Our results in Tables 4 and 5 largely hold when we use OLS regression with the count or log of one plus the count of the number of financial tweets as the dependent variable, though we find weaker support for financial tweet counts reacting to positive news when we use earnings surprise to measure the materiality following Table 5.[18] Lastly, to directly model the counts, we implemented Poisson pseudo maximum likelihood (PPML) regression, using robust standard errors and high-dimensional fixed effects, as in Correia et al. (2020). Our results for Tables 4 and 5 still show the expected U-shaped relation between event materiality and financial tweets.[19]

## 8 | CONCLUSION

This paper examines whether firms symmetrically disseminate financial information around higher-materiality events regardless of the events' direction (positive or negative). Using a large sample of tweets by S&P 1500 firms and an ML approach to classify financial tweets, we find that firms tweet financial information around higher-materiality earnings announcements and 10-K/10-Q filings for both positive and negative events. This result contrasts with those of prior studies using earlier samples, which find either no effect or greater dissemination of good news than bad news on social media. To explore the potential reasons for this difference, we examine

[18]We fully replicate Table 4 using OLS regression with raw or logged counts of financial tweets. Replicating Table 5, we continue to find a U-shaped relation in all columns where it was documented, except for Column 3. However, the main effect of |*Earnings Surprise*| drops significance in Column 3 for both specifications, and the interaction term drops significance in Column 8 for raw counts.

[19]Using PPML regression, we fully replicate Table 4. Replicating Table 5, we continue to find a U-shaped relation in all columns where it was documented, except for Column 3. However, |*Earnings Surprise*| drops significance in Column 3, and the interaction term drops significance in Column 8.

the time trend of firms' tweeting activities. Our results show that although firms react to both extreme good and bad news in the earlier years of our sample period, this pattern is more pronounced in later years. In addition, the magnitude of the effect is larger for later adopters, that is, the firms that joined Twitter during our sample period. These analyses suggest that at least some differences are partially due to sample changes and changes in firms' strategic behavior over time. Furthermore, we examine the intraday pattern of financial tweet disclosure around earnings announcements and 10-K/10-Q filings. We document that on average, financial tweets in the 3 h after earnings announcements do not show the U-shaped relation. On the other hand, financial tweets are frequently posted in the 3 h both before and after 10-K/10-Q filings, and the financial tweets after the filings show the U-shaped relation.

To parse out financial information on Twitter, we introduce a new classification methodology called Twitter-LDA, an unsupervised ML method. We demonstrate that this Twitter-LDA methodology determines a more accurate set of financial information on Twitter than a dictionary approach, and that the increased precision of our Twitter-LDA-based measure is integral to determining the U-shaped relation between news materiality and dissemination in the context of earnings announcements. Our study thus complements early studies on firms' behavior on Twitter by revising our existing belief that managers avoid disseminating bad news. One limitation of our study, however, is that we focus on the strategic behavior of whether firms tweet around accounting events. Social media dissemination is a rich setting with many qualitative characteristics such as the tone and narrative framing of the message; thus many other aspects of dissemination on social media are likely influenced by strategic incentives. Future research may provide interesting insights on these other characteristics.

## ACKNOWLEDGMENTS

We gratefully acknowledge the helpful comments of two anonymous reviewers, Qiang Cheng, Mark Clatworthy, Richard Frankel, Weili Ge, Yu Hou, Xiumin Martin, James Ohlson, Joseph Piotroski, Eddie Riedl, Michael Smith, Phillip Stocken, Mike Welker, T. J. Wong, the seminar participants of Boston University, Brock University, Peking University, Queen's University, University of California at Riverside, University of Toronto, Washington University, and the participants at the SMU-NTU-NUS Accounting Research Conference, CUHK Textual Analysis Conference, CMU Accounting Conference, EAA Annual Congress, Hawaii Accounting Research Conference (HARC), and the *Journal of Accounting, Auditing, and Finance* (JAAF) Conference. This research was supported by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant and the Social Sciences and Humanities Research Council in Canada. Lu acknowledges support from the McCutcheon Professorship in International Business at the University of Toronto.

## DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available in public databases as specified in the paper.

## ORCID

*Richard M. Crowley* https://orcid.org/0000-0001-8488-2351

## REFERENCES

Anilowski, C., Feng, M., & Skinner, D. J. (2007). Does earnings guidance affect market returns? The nature and information content of aggregate earnings guidance. *Journal of Accounting and Economics*, *44*(1–2), 36–63.

Bao, Y., & Datta, A. (2014). Simultaneously discovering and quantifying risk types from textual risk disclosures. *Management Science*, *60*(6), 1371–1391.

Bartov, E., Faurel, L., & Mohanram, P. (2018). Can Twitter help predict firm-level earnings and stock returns? *The Accounting Review*, *93*(3), 25–57.

Beyer, A., Cohen, D. A., Lys, T. Z., & Walther, B. R. (2010). The financial reporting environment: Review of the recent literature. *Journal of Accounting and Economics*, *50*(2–3), 296–343.

Blankespoor, E., & deHaan, E. (2020). Strategic disclosure and CEO media visibility. *Journal of Financial Reporting*, *5*(1), 25–50.

Blankespoor, E., Miller, G. S., & White, H. D. (2014). The role of dissemination in market liquidity: Evidence from firms' use of Twitter™. *The Accounting Review*, *89*(1), 79–112.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, *3*, 993–1022.

Brown, N., Crowley, R. M., & Elliott, W. B. (2020). What are you saying? Using topic to detect financial misreporting. *Journal of Accounting Research*, *58*(1), 237–291.

Bushee, B. J., Matsumoto, D. A., & Miller, G. S. (2003). Open versus closed conference calls: The determinants and effects of broadening access to disclosure. *Journal of Accounting and Economics*, *34*(1–3), 148–180.

Correia, S., Guimarães, P., & Zylkin, T. (2020). Fast Poisson estimation with high-dimensional fixed effects. *The Stata Journal*, *20*(1), 95–115.

Crowley, R. M. (2016). *Market reaction to shifting of voluntary disclosures across disclosure channels*. Working paper. Singapore Management University.

Curtis, A., Richardson, V., & Schmardebeck, R. (2016). Investor attention and the pricing of earnings news. In G. Mitra & Y. Xiang (Eds.), *Handbook of sentiment analysis in finance*. Albury Books.

Dye, R. (1985). Disclosure of nonproprietary information. *Journal of Accounting Research*, *23*(1), 123–145.

Fields, T. D., Lys, T. Z., & Vincent, L. (2001). Empirical research on accounting choice. *Journal of Accounting and Economics*, *31*(1–3), 255–307.

Hassan, H., Hollander, S., van Lent, L., & Tahoun, A. (2024). The global impact of Brexit uncertainty. *The Journal of Finance*, *79*(1), 413–458.

Healy, P. M., & Palepu, K. G. (2001). Information asymmetry, corporate disclosure, and the capital markets: A review of the empirical disclosure literature. *Journal of Accounting and Economics*, *31*(1–3), 405–440.

Hoberg, G., & Lewis, C. (2017). Do fraudulent firms produce abnormal disclosure? *Journal of Corporate Finance*, *43*, 58–85.

Hobson, J., Mayew, W. J., & Venkatachalam, M. (2012). Analyzing speech to detect financial misreporting. *Journal of Accounting Research*, *50*(2), 349–392.

Hollander, S., Pronk, M., & Roelofsen, E. R. (2010). Does silence speak? An empirical analysis of disclosure choices during conference calls. *Journal of Accounting Research*, *48*(3), 531–563.

Hummel, P., Morgan, J., & Stocken, P. C. (2021). *A model of voluntary disclosure*. Working paper. Dartmouth College and University of California.

Jung, M., Naughton, J., Tahoun, A., & Wang, C. (2018). Do firms strategically disseminate? Evidence from corporate use of social media. *The Accounting Review*, *93*(4), 225–252.

Kothari, S. P., Shu, S., & Wysocki, P. D. (2009). Do managers withhold bad news? *Journal of Accounting Research*, *47*(1), 241–276.

Kumar, V., Bhaskaran, V., Mirchandani, R., & Shah, M. (2013). Creating a measurable social media marketing strategy: Increasing the value and ROI of intangibles and tangibles for Hokey Pokey. *Marketing Science*, *32*(2), 194–212.

Larcker, D. F., & Zakolyukina, A. (2012). Detecting deceptive discussions in conference calls. *Journal of Accounting Research*, *50*(2), 495–540.

Lee, L. F., Hutton, A. P., & Shu, S. (2015). The role of social media in the capital market: Evidence from consumer product recalls. *Journal of Accounting Research*, *53*(2), 367–404.

Leuz, C., & Wysocki, P. D. (2016). The economics of disclosure and financial reporting regulation: Evidence and suggestions for future research. *Journal of Accounting Research*, *54*(2), 525–621.

Mayew, W. J., & Venkatachalam, M. (2012). The power of voice: Managerial affective states and future firm performance. *The Journal of Finance*, *67*(1), 1–43.

Merton, R. C. (1987). A simple model of capital market equilibrium with incomplete information. *The Journal of Finance*, *42*(3), 483–510.

Nekrasov, A., Teoh, S., & Wu, S. (2022). Visuals and attention to earnings news on Twitter. *Review of Accounting Studies*, *27*(4), 1233–1275.

Norton, E., Wang, H., & Ai, C. (2004). Computing interaction effects and standard errors in logit and probit models. *The Stata Journal*, *4*(2), 154–167.

Rogers, J. L., Skinner, D., & Zechman, S. (2016). The role of the media in disseminating insider-trading news. *Review of Accounting Studies*, *21*, 711–739.

SEC. (2013, April 2). *SEC says social media okay for company announcements if investors are alerted* [Press release]. https://www.sec.gov/News/PressRelease/Detail/PressRelease/1365171513574

Skinner, D. (1994). Why firms voluntarily disclosure bad news. *Journal of Accounting Research*, *32*(1), 38–60.

Skinner, D. (1997). Earnings disclosures and stockholder lawsuits. *Journal of Accounting and Economics*, *23*(3), 249–282.

Tang, V. (2018). Wisdom of crowds: Cross-sectional variation in the informativeness of third-party-generated product information on Twitter. *Journal of Accounting Research*, *56*(3), 989–1034.

Toubia, O., & Stephen, A. (2013). Intrinsic vs. image-related utility in social media: Why do people contribute to Twitter? *Marketing Science*, *32*(3), 368–392.

Twitter. (2011, September 8). One hundred million voices. *Twitter Blog*. https://blog.twitter.com/official/en_us/a/2011/one-hundred-million-voices.html

Verrecchia, R. E. (1983). Discretionary disclosure. *Journal of Accounting and Economics*, *5*, 179–194.

Zhao, W. X., Jiang, J., Weng, J., He, J., Lim, E., Yan, H., & Li, X. (2011). Comparing Twitter and traditional media using topic models. In P. Clough, C. Foley, C. Gurrin, G. J. F. Jones, W. Kraaij, H. Lee, & V. Murdoch (Eds.), *Advances in information retrieval: Proceedings of the 33rd European Conference on IR Research, ECIR 2011*, Dublin, Ireland (pp. 338–449). Springer.

## SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

**How to cite this article:** Crowley, R. M., Huang, W., & Lu, H. (2024). Discretionary dissemination on Twitter. *Contemporary Accounting Research*, *41*(4), 2454–2487. https://doi.org/10.1111/1911-3846.12986

## APPENDIX 1: VARIABLE DEFINITIONS

| Variable name | Definition |
|---|---|
| ***Tweets*** | Number of tweets posted by the firm on a given day |
| **Dependent variables** | |
| *Financial Tweets, ML* | Indicator equal to one if at least one of the firm's tweets discusses financial information, as classified by Twitter-LDA, on a given day, and zero otherwise |
| *Financial Tweets, Dict* | Indicator equal to one if at least one of the firm's tweets discusses financial information, as classified by a dictionary approach, on a given day, and zero otherwise |
| *Tweeted* | Indicator equal to one if a firm had joined Twitter and had also posted at least one tweet by December 31 of the given year, and zero otherwise |
| *Financial Tweets Hourly, ML* | Indicator equal to one if at least one of the firm's tweets discusses financial information, as classified by Twitter-LDA, in a given hour, and zero otherwise |
| *Financial Tweets Hourly, Dict* | Indicator equal to one if at least one of the firm's tweets discusses financial information, as classified by a dictionary approach, in a given hour, and zero otherwise |
| **Independent variables** | |
| *Earnings Ann* | Indicator equal to one if an earnings announcement is released on a given trading day, and zero otherwise |
| *Form 10-K, 10-Q* | Indicator equal to one if a 10-K or 10-Q filing is released on a given trading day, and zero otherwise |
| *Non-Earnings 8-K* | Indicator equal to one if an 8-K filing is released on a given trading day when no earnings announcement is released, and zero otherwise |

**APPENDIX 1** (Continued)

| Variable name | Definition |
|---|---|
| *Extreme Negative AR* | Indicator equal to one if the day's abnormal return is in the bottom quintile (below or equal to the 20th percentile) of abnormal returns for the firm throughout our sample period |
| *Non-extreme AR* | Indicator equal to one if the day's abnormal return is in the middle three quintiles (above the 20th percentile and below or equal to the 80th percentile) of abnormal return for the firm throughout our sample period |
| *Extreme Positive AR* | Indicator equal to one if the day's abnormal return is in the top quintile (above the 80th percentile) of abnormal return for the firm throughout our sample period |
| [−3, 0) *Hours from Event* | Intraday indicator equal to one for the period 3 h before an event up to the event |
| [0, +3) *Hours from Event* | Intraday indicator equal to one for the period starting at the time of the event up to 3 h after the event |
| **Replication variables** | |
| *Missed Earnings* | Indicator equal to one if the earnings surprise (I/B/E/S) is negative |
| $\lvert \textit{Earnings Surprise} \rvert$ | Absolute value of earnings surprise (from I/B/E/S) |
| $\lvert \textit{Earnings Surprise} \rvert_{\perp}$ | Absolute value of earnings surprise (from I/B/E/S), minus the average value of absolute earnings surprise in the sample |
| *Negative AR* | Indicator equal to one if the abnormal return is negative |
| $\lvert AR \rvert$ | Absolute value of the abnormal return |
| **Control variables** | |
| *Verified* | Indicator equal to one if the firm's Twitter account has been verified, and zero otherwise |
| *Followers* | Number of Twitter followers the firm's Twitter account has |
| *Friends* | Number of accounts that the firm's Twitter account is following |
| *Recent Tweets* | Number of tweets in the 5 trading days leading up to the current day |
| *Total Tweets* | Total number of tweets the firm posted from the time it joined Twitter up to the given trading day |
| *Size* | Natural logarithm of firm's total assets (Compustat: *atq*) |
| *ROA* | Firm's return on assets calculated as net income (Compustat: *niq*) divided by total assets (Compustat: *atq*) |
| *MB* | Market to book ratio, calculated as shares outstanding (CRSP: *shrout*) times shares price (CRSP: *prc*) divided by total assets (Compustat: *at*) |
| *Debt* | Most recent annual long-term debt (Compustat: *ltq*) divided by most recent annual long-term assets (Compustat: *atq*) |
| *Growth* | Quarterly revenue growth calculated as revenue (Compustat: *revtq*) divided by its four-quarter lag |
| *Advertising* | Most recent annual advertising expense (Compustat: *xadv*) |
| *Return Volatility* | Firm's stock return volatility over the past 3 months (63 trading days) |
| *Firm Age* | Firm age based on the date the firm was founded (Compustat *Names* file: *year*1) |
| *Inst Ownership* | Percentage of shares held by institutional ownership (Refinitiv Institutional (13f) Holdings—S34) |
| *log*(*# Analysts*) | Natural log of one plus the number of unique analysts releasing an EPS estimate for the firm in the given quarter (from I/B/E/S) |

(Continues)

**APPENDIX 1** (Continued)

| Variable name | Definition |
|---|---|
| $log(\#\ Analysts)_{\perp}$ | Natural log of one plus the number of unique analysts releasing an EPS estimate for the firm in the given quarter (from I/B/E/S), orthogonalized to firm size as a residual |
| *CEO Age* | Age of the firm's CEO (ExecuComp: *AGE*) |

## APPENDIX 2: TWITTER TOPICS AND EXAMPLES

Each of the subcategories below is comprised of one or more of the 60 topics from the Twitter-LDA algorithm. We manually mapped each of these 60 topics to the aggregations used in the paper: financial, non-financial business, and other. If 40% of a tweet is classified as part of the financial topic, 30% as part of one marketing topic, and 30% as part of one other topic, it will be categorized as a financial tweet, as its most prevalent topic is financial. The first part of the table presents the top 20 words from each subcategory. For subcategories with multiple topics included (indicated in parentheses), we present the words with the highest aggregate weight across all topics within the subcategory. The second part of the table presents representative tweets from the financial category, with three tweets from extreme positive, extreme negative, and non-extreme events. Representative tweets from other subcategories are presented in Appendix S1.

| Subcategory | Top 20 words |
|---|---|
| Financial | market, growth, markets, trading, earnings, global, report, quarter, results, energy, cboe, year, today, week, investors, options, vix, economic, outlook, read |
| Non-financial business | |
| Support (5) | dm, store, customer, team, flight, send, number, hear, feedback, claim, baggage, make, share, relations, confirmation, care, location, switch, socialmedia, apologies |
| Marketing (24) | pass, free, enjoy, shipping, heres, life, love, time, #apple, shop, deals, todays, enter, win, scholarship, check, cr, wed, great, #sharacoke |
| Conference (5) | booth, join, today, #iot, learn, great, live, week, register, stop, power, i, event, design, day, supporting, autograph, video, cancer, tomorrow |
| Other business (8) | #jobs, dm, email, #job, hear, send, contact, hiring, working, details, job, info, store, follow, check, #ljn, number, wed, address, health |
| Other (17) | stay, travelers, dont, rating, order, joe, tweet, collection, enjoy, book, reviews, compare, plan, trip, millions, opinions, perfect, easy, edt, hear |

| Twitter-LDA subcategory | Representative tweets |
|---|---|
| Financial [extreme positive] | • Today we announced our first quarter 2018 earnings. Read more on our performance here: {link}<br>• Summary of $MWW Q1 2013 financial results: E.P.S. was $0.08 . . . EBITDA was $38 M . . .<br>• CEO John Hess: "Our portfolio is on track to generate industry leading cash flow growth and increasing returns to shareholders." See release, including cautionary statement: {link} |

**APPENDIX 2** (Continued)

| Twitter-LDA subcategory | Representative tweets |
|---|---|
| Financial [non-extreme] | • Chevron reports second quarter 2019 earnings $CVX #stocks {link}<br>• Raytheon reports solid first quarter results: {link}<br>• DuPont CEO: We delivered exceptional full-year results despite significant market headwinds late in the year {link}$DD |
| Financial [extreme negative] | • Southern Company reports second quarter earnings {link}<br>• We believe 2014 will mark the turnaround of $STJ #neuromodulation franchise as a growth driver<br>• Hershey's Gold and our Team USA programming are driving category news and growth at retail. {link} $HSY #HSYEarnings |