Essays on Corporate Finance and Risk

by

Bo Sun

(Under the Direction of Jie He and Daniel Rettl)

Abstract

This dissertation is composed of two essays that explore the literature of corporate finance and risk. The first chapter is about expectation management in corporate finance. Expectation management, e.g., via guiding analyst forecasts downwards to achieve positive earnings surprises, is a common practice adopted by corporate management. This paper investigates Twitter as a mechanism for expectation management. CEOs have been increasingly using Twitter to provide information, either about the firm or their personal lives. Using a hand-collected sample of CEO Twitter usernames, I quantify the sentiment of CEOs' Tweets prior to their firms' earnings announcements. I find that tweet sentiment and abnormal stock returns on tweet days are positively correlated, suggesting that the market incorporates the information contained in CEOs' tweets. Furthermore, negative pre-earnings tweet sentiment is followed by downward analyst earnings forecast revisions and positively predicts the likelihood of meeting earnings expectations. The second chapter is about risk management, a core competence of financial institutions and techniques for quantifying market risk that are central to the process of managing risk. Value-at-risk (VaR) and expected shortfall (also called conditional VaR) are two mostly widely used measures of financial risk. However, both measures have some drawbacks. For example, VaR ignores losses beyond a designated threshold and also fails to satisfy mathematical principles characterizing coherent risk measures while expected shortfall fails to have the elicitability criterion deemed essential to backtesting. These deficiencies have motivated interest in other risk measures. Expectile, which was introduced in the context of linear regression, has been revealed as a reasonable risk measure to offset the weaknesses of VaR and expected shortfall. In this paper, using special characterization of the expectile as the minimizer of a suitable discrepancy function, we propose a method to construct a coherent posterior distribution for the expectile and thus provide the full picture of the expectile of the distribution of financial losses through combining the information from both the data and its prior belief. The asymptotic consistency of the posterior is established to support its validity in practice. Some numerical examples are provided to illustrate our method.

INDEX WORDS: [expectation management, earnings surprises, social media, risk measure, expectile, value-at-risk, expected shortfall, posterior distribution]

Essays on Corporate Finance and Risk

by

Bo Sun

B.S., Clemson University, 2017

A Dissertation Submitted to the Graduate Faculty of the
University of Georgia in Partial Fulfillment of the Requirements for the
Degree

Doctor of Philosophy

Athens, Georgia

2022

Essays on Corporate Finance and Risk

by

Bo Sun

Major Professor: Jie He
Daniel Rettl

Committee: Kayla Freeman
Paula Suh
Shuyang Bai

Electronic Version Approved:

Ron Walcott
Vice Provost for Graduate Education and Dean of the Graduate School
The University of Georgia
August 2022

Dedication

To my parents, friends, and family for their support and encouragement through this adventure.

Acknowledgments

I would like to express my deepest gratitude to my dissertation chair, Dr. Jie He. Without his support, guidance, and suggestions, this dissertation would not have been possible. I would like to thank Dr. Daniel Rettl, my co-chair, for all his support throughout my time here in the PhD program. I would also like to thank my committee members, Dr. Kayla Freeman, Dr. Paula Suh, and Dr. Shuyang Bai, for all their suggestions and advice.

I would also like to thank Dr. Jeffry Netter and Dr. Annette Poulsen for their support and guidance. Finally, I would like to thank all my fellow PhD students for their friendship, encouragement, and advice during my time at the University of Georgia.

Contents

List of Figures

List of Tables

Chapter 1

Managing Expectations Through Social Media? CEO Tweets Prior to Earnings Announcements

1.1 Introduction

Managing expectations downwards is a strategy that is used in many applications to help achieve positive surprises. For example, service firms commit to longer delivery time when the actual time is shorter. A car mechanic might tell customers that it will take a week for their car to be fixed, but in reality it only takes five days. When the mechanic gives back the car earlier, a positive surprise is achieved. In corporate finance, firms also use expectation management to make it easier to achieve their goals and attain positive surprises. Matsumoto (2002) analyzes the incentives for managers to engage in expectation management. She finds that firms with transient institutional investors face more pressure to beat earnings forecasts, so managers have incentive to drive analysts' expectations downwards. In addition, in mergers and acquisitions, bidders manage down analyst earnings forecasts prior to earnings releases to increase their own stock prices to save on acquisition costs (He, Liu, Netter, and Shu (2020)). However, the literature on the mechanisms through which managers use to manage expectations is scant and can be extended further, especially by analyzing modern technologies adopted by managers to disclose information.

Starting in the early 2010s, social media emerged as the medium for information dissemination. In financial markets, firms increasingly use social media to disclose financial information (Blankespoor, Miller, and White (2014)), and investors have been retrieving firm information from social media (Brunswick (2014)). Before the passage of Regulation Fair Disclosure, managers can selectively disclose or communicate information to groups of individuals about the firm. However, after the passage, managers need to make public disclosure of that information. Thus, It is interesting to see if managers turn to social media, a public channel, to communicate with analysts and investors, such as managing expectations. To the best of my knowledge, there is no existing research on social media's role in earnings expectation management.

The main benefit of expectation management is to achieve positive earnings surprises on earnings announcement dates, which are salient events that attract a lot of market attention and drive stock valuation. However, it is not costless. The costs of expectation management by utilizing social media are similar to those of traditional methods, such as management forecasts and conference calls. There could be a huge negative stock price reaction to bad news and downward-revised forecasts when these are released to the market, which may not be recovered in the future even considering the positive market reaction on earnings announcement dates. In addition, managers may face damages to their reputation if they are found to manage expectations. Moreover, managers might face legal consequences in the most extreme cases. In the case of social media, litigation risk could be a real and costly consequence. For example, in August 2018, Elon Musk made comments about taking Tesla private on Twitter, which were groundless speculations. As a result, the SEC charged him with civil securities fraud because he never planned to take Tesla private. Also, it seems that using Twitter incurs less transaction costs than other forms of expectation management: Twitter accounts are free and widely accessible for all users, so CEOs can readily and easily use Twitter.

In this paper, I investigate CEO Twitter accounts as a mechanism for CEOs to manage earnings expectations. First, I test if CEOs' Twitter accounts are more active during the month prior to earnings announcement. Second, I analyze if the information content in tweets made prior to earnings announcement is immediately incorporated into stock prices. Lastly, I explore if pre-earnings tweet sentiment affects the likelihood of the firm meeting earnings expectations.

Using a hand-collected sample of CEO twitter usernames, I scrape all tweets made by CEOs from 2006 to 2019. Since Twitter was not very popular in the early years of its creation, I use the sample from 2013 to 2019 for empirical analysis. For each tweet, I use the Python package, Valence Aware Dictionary and sEntiment Reasoner (VADER), to assign it a sentiment score, ranging from -1 to 1, where -1 is the most negative and 1 is the most positive. VADER uses a combination of sentiment lexicons to label text as positive or negative, and it also takes into account the intensity of sentiments . In addition, I classify each tweet as firm-related or non-firm-related using a linear stochastic gradient descent (SGD) classifier algorithm, which adopts a logistic regression.

I find that CEOs tweet 14% more during the month prior to earnings announcement, compared to the month of and the month after earnings announcement. In addition, the sentiment in tweets made prior to earnings announcements is positively associated with tweet day CARs. Furthermore, I find that negative pre-earnings sentiment increases the likelihood of the firm meeting or beating earnings expectations. This finding is consistent with the view that CEOs are using Twitter as a mechanism to manage earnings expectations and ensure positive earnings surprises. Finally, I find that negative pre-earnings tweet sentiment is followed by downward analyst earnings forecast revisions, suggesting that CEOs achieve such expectation management through successfully lowering analyst expectations.

This paper contributes to the prior literature in several ways. First, the paper extends the literature on expectation management by providing additional evidence that CEOs manage earnings expectations to achieve positive surprises (Matsumoto (2002), Bartov et al (2003)). Second, this paper shows that social media, more specifically, Twitter, is a modern yet unexplored mechanism that can be used to manage expectations. I find that prior to earnings announcements, CEOs' Twitter accounts are more active and informative, and tweet sentiment is associated with the likelihood of beating earnings expectations and analyst revisions. This paper highlights the growing importance of alternative ways for managers to communicate and disclose information, and investors should understand the implications of such alternative methods.

Finally, this paper contributes to the growing literature on social media. Prior studies focus on the impact of social media on financial markets and stock returns (e.g., Chen et al (2014), Antweiler and Frank (2005)). Other papers such as Gao (2019) investigate the information content of CEO Twitter accounts, but do not explore Twitter as a mechanism for expectation management. To the best of my knowledge, this is the first paper to use Twitter as a channel for CEOs to disclose information for the purpose of expectation management.

The rest of the paper is structured as follows. Section 2 provides the prior literature and hypotheses. Section 3 details the data and methodology. Section 4 and 5 describe the empirical tests. Section 6 provides cross-sectional analysis, and Section 7 contains the robustness tests. Lastly, Section 8 concludes.

1.2 Related Literature and Hypotheses

1.2.1 Previous Literature

Prior studies find that the market punishes firms for reporting negative surprises. Skinner and Sloan (2002) find that firms that announce negative earnings surprises exhibit an asymmetrically large negative price response. Similarly, Brown and Caylor (2005) provide evidence that the market penalizes firms more for missing analysts' forecasts than for reporting earnings decreases. On the other hand, firms are rewarded for meeting or beating earnings expectations. Bartov et al (2002) find that firms that meet or beat analysts' earnings forecasts have higher returns for the next quarter than firms falling short of the expectation. Furthermore, the cost of breaking consecutive earnings announcements where the firm meets the expectation is high (Kross et al (2011)).

Therefore, managers have incentives to meet or beat analysts' earnings expectations. Matsumoto (2002) analyzes the incentives for managers to engage in expectation management. He reveals that firms with transient institutional investors face more pressure to beat earnings forecasts, so managers have incentive to drive analysts' expectations downwards. In addition, in mergers and acquisitions, bidders manage down analyst earnings forecasts prior to earnings releases to increase their own stock prices to save on acquisition costs (He, Liu, Netter, and Shu (2020)). Overall, these results suggest that managers are not passive observers, but rather try to guide analyst expectations downwards.

The expectation management literature is consistent with the "preparing" the equity market literature. Chemmanur and Tian (2011) find that it is beneficial for some firms to "prepare" the market for adverse corporate events. They develop a theoretical model to generate several predictions, such as less transparent firms are more likely to prepare the market. These findings highlight the importance of firms voluntarily disclosing information.

Since managers are incentivized to engage in expectations management, prior studies have investigated the mechanisms in which managers can provide these pessimistic guidance. In October 2000, Regulation Fair Disclosure, or Reg FD, prevented companies from privately communicating information to selected analysts or investors, without revealing it publicly. Therefore, in the post Reg FD period, managers turn to voluntary disclosures, such as conference calls, to manage analyst expectations (Baik and Nam (2009)). In this paper, I investigate social media as a mechanism for managers to engage in expectation management.

There has been some literature on the impact of social media on financial markets, ranging from traditional social media, such as Facebook, to messaging boards such as Yahoo! and Seeking Alpha (Chen et al (2014), Antweiler and Frank (2005), Giannini, Irvine, and Shu (2019)). In this paper, I utilize CEO Twitter accounts as my disclosure channel. First, Twitter provides information that may not be present in traditional disclosures, such as earnings announcements and M&A announcements. A CEO's or firm's Twitter account provides day-to-day managerial information, personal thoughts and mood by the CEO, etc. Second, Twitter is different from other social medias like Facebook, because Twitter increases the repetition of information through multiple tweets and retweets of selected content (Miller and Skinner (2015)).

Gao (2019) also collects personal CEO tweets, and finds that positive sentiment predicts positive future abnormal returns. In addition, she provides evidence that positive sentiment predicts positive earnings surprises because she shows that positive CEO sentiment is a proxy for firm performance. However, the literature on social media as a mechanism for expectation management is lacking. Therefore, this paper's purpose is to examine social media as a possible channel.

1.2.2 Hypothesis Development

Managers are incentivized to guide expectations downwards as earnings announcement date approaches. In this paper, I investigate social media's role in expectation management. Specifically, I examine CEO Twitter accounts as a mechanism for CEOs to drive analyst expectations downwards to increase the likelihood of meeting earnings expectations. If CEO Twitter accounts are involved in expectation management, then it must be a viable communication channel for CEOs to the public.

First, I examine the usage of CEO Twitter accounts. If Twitter is used to guide analyst expectations downwards prior to earnings announcement dates, then I expect that there will be more tweets prior to earnings announcements than at other times. Therefore, I formulate the first hypothesis below:

Hypothesis 1: CEOs tweet more in the month prior to earnings announcements than during or the month after.

After establishing the prevalent usage of CEO Twitter accounts, I investigate the stock market reaction to the information content of the tweets. If pre-earnings tweet sentiment on CEO Twitter accounts guides analyst expectations downwards, then pre-earnings tweets should contain valuable information for the market. I generate the hypothesis below:

Hypothesis 2: Tweet sentiment should be positively correlated with tweet date CARs.

Tweets made on CEO accounts contain a variety of information, such as firm-related or non-firm related. Firm-related information contains important news for the firm, such as new products or firm events. Even if the news is not new information, mentioning old news still has a similar effect (Tetlock (2011)). On the other hand, non-firm-related tweets provide information about the CEO's day-to-day personal life, such as family events and mood. Therefore, it is interpreted by investors as proxies for how well the firm is doing. In either instance, I expect the information on Twitter predicts tweet day CARs.

Lastly, I analyze if CEOs use their Twitter accounts as a way to manage analysts' earnings expectations. I expect that pre-earnings sentiment is associated with an increased likelihood of beating earnings estimates. Therefore, I propose the following hypothesis:

Hypothesis 3: Firms with CEOs who tweet negative sentiment prior to earnings announcements are more likely to meet/beat earnings targets.

In the accounting literature, various papers have shown that managers take actions to avoid negative earnings surprises (Brown and Caylor (2005) and Matsumoto (2002)). The sentiment on CEO Twitter accounts should be a great predictor of the firm meeting earnings expectations. Specifically, negative sentiment increases the likelihood of the firm meeting earnings expectations, which is driven by the analysts revising their estimates lower. Therefore, I expect negative pre-earnings sentiment increases the likelihood of analysts revising their estimates downwards.

Taken altogether, this paper provides evidence that CEOs use their personal Twitter accounts to manage analyst expectations. In particular, Twitter is used as a mechanism, along with traditional methods such as conference calls, to lower earnings expectations.

1.3 Data and Methodology

To study the importance of alternative information in social media, I construct the sample by manually identifying the personal Twitter accounts of CEOs. The Twitter usernames are hand-collected from a sample of 2,726 current CEO names from Execucomp from 2006 to 2019, since Twitter was created in 2006. I collect 187 Twitter usernames from the list of 2,726 CEO names, so a vast majority of current CEOs listed in Execucomp had no Twitter accounts. The usernames are found through Google, by searching the CEO names and their companies, to make sure the Twitter username corresponds to the right person. Using the list of Twitter usernames, I use a Python package called Twitterscraper to scrape all tweets made by each CEO on their personal Twitter accounts. I only include tweets made during the CEO's tenure, so I can analyze their personal Twitter accounts when they are working at the firm.

I collect a total number of 156,883 tweets from the 187 CEO twitter accounts above. Out of the 187 Twitter accounts analyzed, 151 Twitter accounts have tweeted at least once during the time when the CEO is actively employed in the firm. For each tweet, I obtain the following data: the text of the tweet, number of likes, number of retweets, timestamp, hashtags used, replies, and if the tweet contains any media (e.g. pictures, videos).

Figure 1 shows the number of CEO tweets published per year, starting from 2006 to 2019. During the time when Twitter first launched, there are not many CEOs who tweeted, or even had an account. However, during the 2010s, Twitter started to rise in popularity, and many CEOs started to register for accounts. In 2011, the total number of tweets by CEOs increased to 5000. Four years later, the total number of tweets exponentially increased, to as high as 25,000 tweets per year. CEOs have exponentially increased their presence on Twitter to a point where investors cannot ignore their social media interactions.

Since CEOs are mostly inactive on Twitter from 2006 to 2012, I exclude these years in my sample. Furthermore, the analysis focuses on pre-earnings tweets, so I do not include tweets that are not around earnings announcements. Using the 156,883 tweets acquired, I aggregate all tweets each day to provide tweet-day information. Tweets made after 4pm are pushed back to the next trading day. All tweet characteristics, such as number of likes, replies, and retweets, are also aggregated by the day. This empirical setup was chosen, so that I can run, for example, event study analysis in later sections. Moreover, I only include tweets made 30 days prior to earnings announcements.

For each tweet, I use the Valence Aware Dictionary and sEntiment Reasoner, or VADER, to assign a sentiment score. VADER is a lexicon and rule-based sentiment analysis tool that is specifically attuned to sentiments expressed in social media. It handles various language criteria, such as negations, contractions, punctuation, and slang. Figure 2 shows examples of tweets and their sentiment scores. The tweets shown consists of one example of each category: Negative firm-related, Positive firm-related, negative non-firm-related, and positive non-firm-related tweets. For example, Panel A shows the negative firm-related tweet by Devin Wenig who is the CEO of Ebay. He tweets that Ebay will be impacted by the newly imposed internet sales taxes. The negative words are imposed and lower, and the total sentiment score for this tweet is -0.1779.

I obtain firm stock price data from the Center for Research in Security Prices (CRSP), and accounting data from the Compustat database. I also collect analyst forecast data and earnings announcement data from the Institutional Brokers' Estimate System (I/B/E/S). I winsorize all continuous variables at the 1% and 99% levels to exclude outliers.

Furthermore, I acquire important news and events from Capital IQ Key Developments. I drop all tweet days where there are any important news and events listed to alleviate concerns that tweet days are affected by other events, such as executive changes, M&A rumors, or changes in corporate guidance, etc. This is done to avoid any days that are hugely affected by other factors. The final sample size is 8,849 tweet days from 2013 to 2019.

Table 1 shows the summary statistics of the sample used in the main empirical analysis. Panel A shows the daily tweet variables which is constructed by aggregating all the tweets made in one day. All the variables are winsorized at the 1% and 99% levels to account for outliers. Notice that Sentiment Score ranges from -1 to 1, where -1 is the most negative and 1 is the most positive. Around 20% of the tweet days contain negative language, hence they have a Negative Score. Most tweet days have positive sentiment, with a median positive sentiment of 0.40. Likes, replies, and retweets are skewed to the right, highlighting that the right tail extremes contain tweets with vastly greater amount of likes, replies, and retweets. In Panel B, firms in our sample beat earnings on average 74% of the time, and 49% of analyst revise their estimates downwards in the last month prior to earnings. Firm controls are shown in Panel C.

1.4 Tweet Sentiment and Stock Returns

1.4.1 Tweeting Frequency Near Earnings

First, I examine the usage of CEO Twitter accounts. If Twitter is used to guide analyst expectations downwards prior to earnings announcement dates, then I expect that there will be more tweets prior to earnings announcements than at other times. I consider the following model:

$$ln(1 + Number\ of\ Tweets)_{i,t} = \alpha_i + \alpha_t + \sum_{n=-1}^{1} \beta_n Earnings_{t+n} + \gamma X_{i,t-1} + \epsilon_{it},$$

where $Earnings_{t-1}$ is an indicator variable that equals one if tweets are one month before earnings. $X_{i,t-1}$ represents the corresponding firm controls which are calculated one year prior to tweet date. CEO and Year Fixed Effects are added in the specification.

Table 2 presents the results. The coefficient on $Earnings_{t-1}$ is positive and statistically significant. In order to show that CEOs tweet more in the month prior to earnings announcement, the coefficient on $Earnings_{t-1}$ is compared to $Earnings_t$ and $Earnings_{t+1}$, using the F-test. The results of the F-test are shown at the bottom of Table 2. The F-values of 9.18 and 27.46 shows that there are more tweets made one month prior to earnings announcement than the month during and the month after, respectively. It suggests that CEOs are more active on Twitter the month prior to earnings, which is consistent with CEOs using Twitter as a mechanism to manage earnings expectations.

1.4.2 Sentiment Score and Tweet Day CARs

Next, I show that the information contained in tweets is incorporated into the market. Consistent with prior literature, I expect a positive association between sentiment score and CAR around tweet date (Gao (2019), Chen et al (2014), Antweiler and Frank (2005)). I consider the following regression model:

$$\text{Tweet Day CAR}(0, +1)_{i,t} = \alpha_i + \alpha_t + \beta\ \text{Sentiment Score}_{i,t} + \delta Z_{i,t} + \gamma X_{i,t-1} + \epsilon_{it}$$

where $TweetDayCAR(0, +1)_{i,t}$ is the cumulative abnormal return around tweet date, and Sentiment $\text{Score}_{i,t}$ is the mean VADER sentiment score for tweets made for CEO i in day t. $\delta Z_{i,t}$ is the tweet characteristics, such as log number of likes, replies, and retweets. $X_{i,t-1}$ represents the corresponding firm controls which are calculated one year prior to tweet date. CEO and Year Fixed Effects are added in the specification.

I expect the coefficient on *Sentiment Score* to be positive, as positive tweets predicts positive CARs, and negative tweets predict negative CARs. Table 3 shows the regression results of tweet day CARs on daily tweet sentiment scores. Panel A shows the results for using Tweet Day CARs on the date of tweet and the day after. Panel B shows the results for a -1 to 1 window. The coefficient on *Sentiment Score* is positive and significant for all specifications, suggesting that tweet sentiment is positively correlated with tweet day CARs. The results are similar when using Fama-French 3 factor model in columns 5 to 8. Taking all together, the results in Table 3 show that the market incorporates the information provided in tweets, suggesting the information is important.

1.4.3 Negative and Positive Scores

Since the hypothesis in this paper focuses on negative sentiment, I decompose *Sentiment Score* into two categories: *Negative Score* and *Positive Score*. *Negative Score* equals the absolute value of the sentiment score if the sentiment score is negative, and zero otherwise. Therefore, it ranges from 0 to 1, with zero being the least negative, and one being the most negative. Similarly, *Positive Score* equals the value of the sentiment score if the sentiment score is positive, and zero otherwise. Therefore, it ranges from 0 to 1, with 0 being the least positive, and 1 being the most positive. The regression is shown below:

$$\begin{aligned}\text{Tweet Day CAR}(0,+1)_{i,t} &= \alpha_i + \alpha_t + \beta_1 \text{ Negative Score}_{i,t} \\ &+ \beta_2 \text{ Positive Score}_{i,t} \\ &+ \delta Z_{i,t} + \gamma X_{i,t-1} + \epsilon_{it}\end{aligned}$$

where $TweetDayCAR(0,+1)_{i,t}$ is the cumulative abnormal return on tweet date. *Negative Score*$_{i,t}$ is the mean VADER negative sentiment score for tweets made for CEO i in day t, and *Positive Score*$_{i,t}$ is the mean VADER positive sentiment score for tweets made for CEO i in day t. $Z_{i,t}$ is the tweet characteristics, such as log number of likes, replies, and retweets. $X_{i,t-1}$ represents the corresponding firm controls which are calculated one year prior to tweet date. CEO and Year Fixed Effects are added in the specification.

The results of the deconstructed sentiment score are presented in Table 4. Panel A shows the results with CARs in a [0,+1] window, and Panel B uses a window of [-1,+1]. The coefficient on *Negative Score* is negative and significant, while the coefficient on *Positive Score* is statistically significant, but small in magnitude. This suggests that CARs after negative sentiment have higher magnitudes than CARs after positive sentiment. Therefore, pessimistic tweets are incorporated into stock price much better than positive tweets. Panel B uses the indicator sentiment variables that equals one if the continuous variables are nonzero, and equals zero, otherwise. A possible explanation is that there is a lot of noise in positive sentiment tweets because the majority of tweets have positive sentiment which includes tweets about not as serious news. However, negative tweets are more likely to contain serious and more important news, which leads to less noise.

1.5 Pre-earnings Tweet Sentiment and Earnings Announcements

1.5.1 Meeting Earnings Expectations

If CEOs use their personal Twitter accounts to manage earnings expectations, then disclosing bad news on Twitter prior to earnings announcements should increase the likelihood of meeting earnings expectations. I analyze if firms meet earnings expectations more often when the CEO tweets out negative pre-earnings sentiment on Twitter. Therefore, I establish the following logistic regression model:

$$\text{Beat}_{i,t} = \alpha_i + \alpha_t + \beta \text{ Monthly Sentiment Score}_{i,t-1} + \delta Z_{i,t} + \gamma X_{i,t-1} + \epsilon_{it},$$

where $Beat_{i,t}$ is an indicator variable that equals one if there firm meets earning expectations and equals zero, otherwise. The latest mean analyst expectations are used to see if the firm meets or does meet expectations. *Sentiment Score*$_{i,t-1}$ is the mean VADER sentiment score for tweets made for CEO i in month $t-1$. $Z_{i,t}$ is the tweet characteristics, such as log number of likes, replies, and retweets. $X_{i,t-1}$ represents the corresponding firm controls which are calculated one year prior to earnings announcement date. CEO and Year Fixed Effects are added in the specification.

Table 5 shows the logit regression results of beating earnings expectations on prior month tweet sentiment. The dependent variable, *Beat Earnings Indicator*, is a dummy variable that equals one if the firms meet the most current earnings expectations, and equals zero, otherwise. Panel A uses the continuous sentiment score, where *Sentiment Score* ranges from -1 to 1. Using the same construction previously, *Sentiment Score* is decomposed into negative and positive scores, which range from 0 to 1. I find that prior month *Negative Score* is positively correlated with beating earnings expectations, suggesting that pre-earnings negative sentiment reduces earnings expectations, leading to a higher chance to beat earnings. Panel B uses indicator variables to denote pre-earnings tweet sentiment, and we find similar results.

1.5.2 Downward Analyst Revisions

If pre-earnings tweet sentiment leads to a higher probability to meet earnings expectations, then it should be positively correlated with analysts revising their estimates downwards. Therefore, I test if analysts revise their estimates downwards in the month prior to earnings, if the CEO tweets out negative sentiment in the second month prior to earnings. Therefore, I build the following logistic model:

$$\begin{aligned}\text{Downward Revision}_{i,t-1} = \alpha_i + \alpha_t + \beta\text{Monthly Sentiment Score}_{i,t-2} \\ + \delta Z_{i,t-2} + \gamma X_{i,t-1} + \epsilon_{it},\end{aligned}$$

where *Downward Revision*$_{i,t-1}$ is an indicator variable that equals one if an analyst revises his/her estimates downwards in the month prior to earnings. *Sentiment Score*$_{i,t-1}$ is the mean VADER sentiment score for tweets made for CEO i in month $t-2$. $Z_{i,t-2}$ is the tweet characteristics, such as log number of likes, replies, and retweets. $X_{i,t-1}$ represents the corresponding firm controls which are calculated one year prior to earnings announcement date. CEO and Year Fixed Effects are added in the specification.

The results of this are shown in Panel A of Table 6. I regress *Downward Revision* on the tweet sentiment. *Downward Revision* is an indicator variable that equals one if an analyst revises his/her EPS estimates downwards in the month prior to earnings. Tweet sentiment variables are aggregated using tweets made between 60 and 30 days prior to earnings announcement. I find that the tweet sentiment in the second month prior to earnings announcement is negatively correlated with downward analyst revisions. Moreover, negative tweet sentiment drives this results and is positively correlated, providing evidence that negativity on Twitter predicts analysts lowering their EPS estimates. Consistent with the hypothesis that firms beat earnings more when there is negative sentiment prior, downward analyst revisions is one channel that this occurs.

Panels B and C of Table 6 show the robustness tests of this analysis. *Upward Revision* is an indicator variable that equals one if an analyst revises his/her EPS estimates upwards in the month prior to earnings. I do not find any statistical significance on the coefficients of sentiment variables. Panel C uses the number of downward analyst revisions as the dependent variable. Using this continuous variable, the results are similar compared to Panel A.

1.6 Cross-Sectional Analysis

1.6.1 Firm-Related vs Non Firm-Related

A CEO's Twitter account may contain firm-related information or personal day-to-day information. Intuitively, these two types of information should have different effects because firm-related information may be more valuable. Firm information disclosed on CEO Tweets may not be present in other sources which makes it very important for analysts and investors to be familiar with. On the other hand, personal day-to-day information is also important because this gives us insight on the CEO's mood. If a CEO knows that negative tweet sentiment can manage analyst expectations downwards, it is rational to think that CEOs may try to induce the same effect with personal tweets.

Using machine learning, I classify each tweet as firm-related and non firm-related. First, I manually classify 1000 CEO tweets as firm-related and non firm-related, by seeing if the tweet mentions anything about the firm, such as firm news, new projects, etc. Then, using 500 of those tweets as the training sample and 500 tweets as the test sample, I fit a logistic regression using stochastic gradient descent to optimize the fit. To be more specific, the algorithm breaks the text of a tweet into tokens, and assigns probabilities of words that appear in firm-related tweets. The model achieves a 92% success rate by testing the fit on our test sample. Then, I split our sample into two sub-samples and repeat the analysis.

First, I test if there is any difference between tweet frequency of firm-related and non firm-related tweets. It is important to note that there are more non firm-related tweets than firm-related tweets. An explanation for this is that CEOs may use their personal Twitter accounts more for day-to-day personal life, unlike the Twitter accounts for firms. In Table 7, I repeat our analysis for the two sub-samples. Panel A shows the results for firm-related tweets, and we could see that there is a 93% increase of firm-related tweets occurring in the one month prior to earnings than during earnings. On the other hand, Panel B shows that there is a 62% increase of non firm-related tweets occurring in the one month prior to earnings than during earnings. Therefore, firm-related tweets increase more for the month prior to earnings than non firm-related tweets. An explanation for this is that CEOs are tweeting out more firm-related news before earnings to try to prepare the market for the firm's earnings announcement. Another reason is that CEOs are often more busy on firm-related work in the month prior to earnings, so they have less time to tweet out personal tweets.

Second, I repeat our analysis to see if there are differences in stock market reaction. I predict that firm-related tweets have a greater impact on tweet date CARs because the importance of the company news outweighs personal news. In Table 8, I find that the estimate of the coefficient on Sentiment Score for firm-related is 0.340%, which is larger than coefficient for non firm-related tweets.

Tables 9 and 10 show the regression results for firms beating analyst expectations and downward analyst revisions. The results are similar compared to the whole sample, and provide evidence that firm-related tweets are stronger in magnitude.

1.6.2 Effectiveness of CEO Twitter Accounts Over Time

In this section, I examine if CEO Twitter Accounts vary in effectiveness depending on the age or time the CEO has spent on Twitter. CEOs join Twitter during various times, either in the early stages of Twitter or recently. It is interesting to see if there are any cross-sectional differences between CEOs joining at different times.

There are two possible competing hypotheses regarding the age or time of the Twitter account. A CEO Twitter account may be more effective the longer it is on the platform because participation increases credibility. The audience believes the information if the CEO builds his or her relationship with the Twitter participants. On the other hand, investors may realize that the CEO is using Twitter to manage expectations. Therefore, their expectations are not reflected by the Twitter sentiment. Specifically, I consider the following logistic regression:

$$
\begin{aligned}
\text{Beat}_{i,t} = \alpha_i + \alpha_t &+ \beta_1 \text{ Monthly Sentiment Score}_{i,t-1} \times \text{Account Age}_{i,t} \\
&+ \beta_2 \text{ Monthly Sentiment Score}_{i,t-1} + \beta_3 \text{ Account Age}_{i,t} \\
&+ \delta Z_{i,t-1} + \gamma X_{i,t-1} + \epsilon_{it}
\end{aligned}
$$

where $Beat_{i,t}$ is an indicator variable that equals one if there firm meets earning expectations and equals zero, otherwise. The latest mean analyst expectations are used to see if the firm meets or does not meet expectations. *Sentiment Score*$_{i,t-1}$ is the mean VADER sentiment score for tweets made for CEO i in month $t-1$. *Account Age* is defined as the number of months from the account's first tweet made. $Z_{i,t}$ is the tweet characteristics, such as log number of likes, replies, and retweets. $X_{i,t-1}$ represents the corresponding firm controls which are calculated one year prior to earnings announcement date. CEO and Year Fixed Effects are added in the specification.

Table 11 presents the results for this test. Similar to before, the coefficients on *Negative* and *Sentiment Score* are statistically significant, consistent with the results that pre-earnings sentiment is associated with the likelihood of meeting earnings expectations. The coefficient on *Account Age* is positive and statistically significant, which implies that CEOs that have been on Twitter longer increases the likelihood of beating earnings. A possible explanation is that investors may see CEOs' tweets as more credible the longer he/she has been on Twitter. Moreover, social media is growing larger, so the audience is increasing. However, the coefficient on the interaction term is not statistically significant, implying that there is no diminished or amplified association between tweet sentiment and meeting earnings expectations, varying on account age. This suggests that CEOs using their personal Twitter accounts to manage expectations is just the beginning, and the importance of social media is not going away anytime soon.

1.7 Robustness Tests

1.7.1 Extreme Tweet Sentiment

In Section 4.2, I have shown that tweet sentiment prior to earnings announcements is associated with tweet day CARs. Negative tweet sentiment predicts negative tweet day CARs, and positive tweet sentiment predicts positive tweet day CARs. One interesting research question is if this correlation is driven by the extreme sentiments per day. On a particular day, an extremely negative tweet may overshadow other tweets and drive the results. Therefore, I repeat the analysis in Section 4.2, and use the maximum values of negative and positive sentiment scores, instead.

Table 12 shows the regression results using the maximum negative and positive scores for each tweet day. The coefficients on *Max Negative* and *Max Positive* are both statistically insignificant across all specifications. I find that the extreme tweet does not drive the results, and does not overwhelm the sentiment from other tweets. One possible explanation for the statistical insignificance is that only using the most extreme tweet sentiment excludes the other tweets made on the same day. Therefore, it removes the sentiment from the various tweets that are excluded which can have valuable information.

1.7.2 Single Tweet Days

In this section, I investigate if the correlation between tweet sentiment and tweet day CARs hold if I only include tweet days in which only one tweet was made.

This analysis is used to see if even the smallest amount of information, just one tweet, is reflected in stock prices. Therefore, I repeat the analysis in Section 4.2 by only including tweet days containing one single tweet.

Table 13 shows the regression results using only tweet days which contain only one tweet made. The results here are similar to that of the whole tweet sample. The coefficient on *Negative* is statistically significant and negative, which is consistent with negative tweets being correlated with negative stock price reactions. However, the coefficient on *Positive* is not statistically significant, which implies that positive sentiment tweets do not have as much power in predicting abnormal stock returns. A possible explanation is that the majority of CEO tweets are positive, so the market does not pay attention to positivity as much as negativity.

1.7.3 Meeting Earnings Expectations: Falsification Test

In this section, I present a falsification test by using the prior quarter's EPS as targets. This empirical setup provides further robustness checks for the correlation of pre-earnings tweet sentiment and meeting earnings estimates. If pre-earnings tweet sentiment is a good predictor of the firm beating earnings expectations, then I expect the increased likelihood is driven by the decreases in analyst forecast estimates made after the tweets.

Table 14 shows the logit regression of *Beat Prior Quarter* on tweet sentiment scores, where *Beat Prior Quarter* is an indicator variable that equals 1 if a firm's earnings per share exceeds the expected earnings per share, using the prior quarter's EPS as the benchmark. In this robustness test, I find that the coefficients on *Negative Score, Positive Score, and Sentiment Score* are all statistically insignificant. The economic magnitude on these coefficients are also small. *Negative Score* has a coefficient of -0.097, *Positive Score* has a coefficient of -0.033, and *Sentiment Score* has a coefficient of -0.026. This provides further support that the increased likelihood of beating earnings expectations is likely driven by the changes in analyst forecast revisions.

1.7.4 Marginally Beating Earnings Expectations

Another explanation is that CEOs manage expectations just enough to beat earnings estimates. Therefore, I expect that when a CEO engages in management expectation, the firm's earnings per share should be just marginally above the managed analyst forecasts. I repeat the analysis in Section 5.1 by using an indicator variable that equals one if the firm meets earnings expectations by 0.01.

Table 15 shows the logit regression results of *Close Beat*, which is an indicator variable that equals 1 if the firm meets earnings expectations and the firm's *Earnings Surprise* does not exceed $0.01, where *Earnings Surprise* is defined as the actual EPS minus the mean analyst estimate in the prior month, divided by price. It equals 0 if the firm misses its earnings target. All firms meeting their earnings expectations with *Earnings Surprise* exceeding $0.01 are excluded from the sample. The results are similar to that of Table 5, which signals that managers use Twitter as a mechanism to guide analyst estimates downwards to marginally beat earnings expectations.

1.8 Conclusion

This paper explores CEO Twitter accounts as a mechanism for CEOs to manage earnings expectations. CEOs tweet the most in the month prior to earnings announcement. I provide evidence that pre-earnings tweet sentiment is positively correlated with CARs on tweet date, suggesting that the market digests the information presented in CEO personal Twitter accounts. Furthermore, I show that negative pre-earnings sentiment increases the likelihood of the firm meeting earnings expectations. Firms with CEOs that are pessimistic prior to earnings are associated with downward analyst revisions, which is consistent with the literature on managers choosing to disclose bad news to avoid any negative earnings surprises.

An implication of this paper is that investors need to pay attention to social media because they contain a vast amount of useful information. On the other hand, CEOs or managers use social media as another source of providing information, in addition to the traditional sources such as news presses and filings. Therefore, lawmakers need to pay attention to the effect of social media on financial markets because it is easier than ever to release and disclose information. In addition, investors need to pay attention to managers using social media to manage expectations because positive surprises may be induced by management.

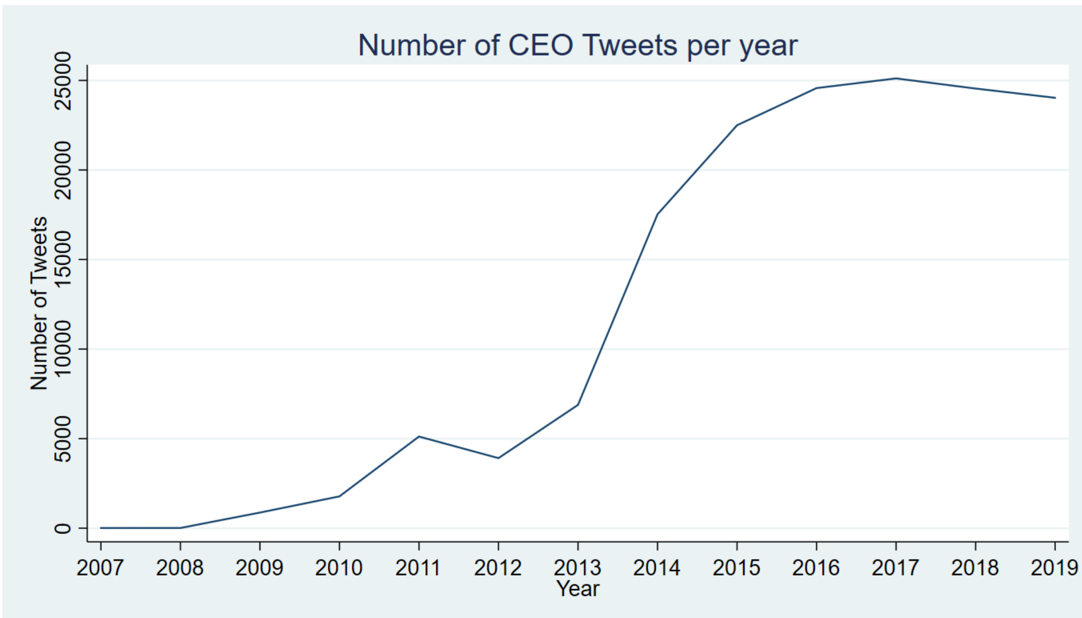

Figure 1.1: Plots the annual CEO tweeting frequency across all firms by year.

Figure 1.2: Examples of CEO Tweets:

Panel A: Negative Firm-Related Tweet

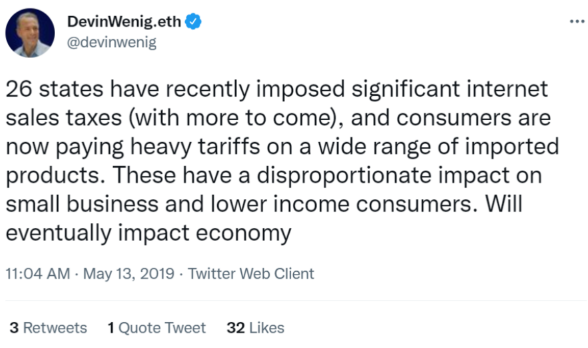
DevinWenig.eth
@devinwenig

26 states have recently imposed significant internet sales taxes (with more to come), and consumers are now paying heavy tariffs on a wide range of imported products. These have a disproportionate impact on small business and lower income consumers. Will eventually impact economy

11:04 AM · May 13, 2019 · Twitter Web Client

3 Retweets 1 Quote Tweet 32 Likes

Sentiment Scores

- Imposed : -0.0772
- Lower: -0.296
- Total tweet: -0.1779

Panel B: Positive Firm-Related Tweet


Lisa Su
@LisaSu

Beautiful racks of @AMD #EPYC servers. Excited to have our lead customers, partners and press join us in Austin today for our launch event!

9:24 AM · Jun 20, 2017 · Twitter Web Client

269 Retweets 28 Quote Tweets 918 Likes

Sentiment Scores

- Beautiful: 0.5994
- Excited: 0.34
- Total tweet: 0.8313

Panel C: Negative Non-Firm-Related Tweet

Sentiment Scores

- Violence : -0.6249
- Failing: -0.5106
- Total tweet: -0.903

Panel D: Positive Non-Firm-Related Tweet

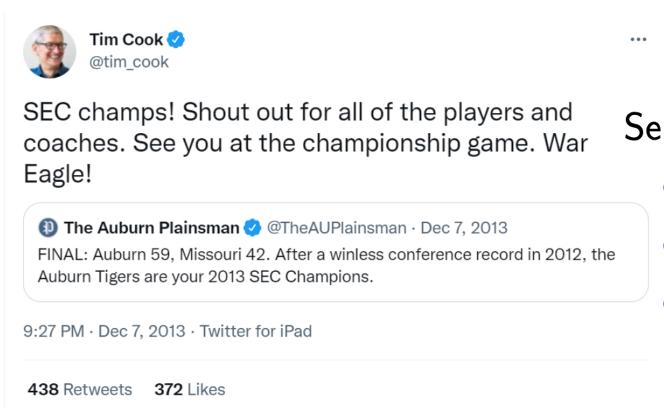

Sentiment Scores

- Champs: 0.4215
- Championship: 0.4404
- Total tweet: 0.9467

Table 1.1: Summary Statistics

This table reports the summary statistics of the variables used in this paper. Panel A shows the aggregated daily tweet summary statistics. Panel B shows the variables used in the earnings expectations section. Panel C shows the summary statistics for firm controls. Definitions of all the variables are reported in Appendix A.

Panel A. Daily Tweet Sample

	Obs	Mean	SD	Min	p25	Median	p75	Max
Tweet Day CAR(0,+1)	8849	−0.01	2.69	−9.18	−1.25	−0.01	1.21	10.12
Tweet Day CAR(-1,+1)	8849	−0.00	3.13	−10.26	−1.52	0.00	1.47	10.71
Tweet Day FF3(0,+1)	8849	0.05	2.63	−8.83	−1.17	0.01	1.20	10.20
Tweet Day FF3(-1,+1)	8849	0.08	3.06	−10.09	−1.39	0.07	1.49	11.01
Sentiment Score	8849	0.38	0.34	−0.93	0.10	0.40	0.66	0.99
Negative Score	8849	0.02	0.10	0.00	0.00	0.00	0.00	0.93
Positive Score	8849	0.40	0.30	0.00	0.10	0.40	0.66	0.99
Likes	8849	282.12	1011.97	0.00	1.00	8.00	46.00	7630.00
Replies	8849	21.42	68.31	0.00	0.00	0.00	3.00	433.00
Retweets	8849	83.33	270.98	0.00	0.00	3.00	18.00	1892.00

Panel B. Variables of Interest

	Obs	Mean	SD	Min	p25	Median	p75	Max
Log(Monthly Tweets)	6588	1.79	1.60	0.00	0.00	1.61	3.00	4.96
Beat Earnings	1174	0.74	0.44	0.00	0.00	1.00	1.00	1.00
Earnings CAR(0,+1)	1137	0.38	8.21	−26.00	−3.58	0.55	4.27	23.32
Earnings Surprise	1137	0.01	0.04	−0.14	0.00	0.01	0.02	0.20
Downrevision	1137	0.49	0.50	0.00	0.00	0.00	1.00	1.00

Panel C. Firm Controls

	Obs	Mean	SD	Min	p25	Median	p75	Max
ROA	8849	0.12	0.11	−0.09	0.07	0.10	0.15	0.75
MB	8849	2.75	2.20	0.80	1.37	1.94	3.44	12.93
Leverage	8849	0.21	0.17	0.00	0.05	0.20	0.33	0.63
Ln(Assets)	8849	8.41	2.06	4.99	6.63	8.36	9.90	12.58
Capex	8849	0.32	0.18	0.03	0.19	0.28	0.42	0.80
RD	8849	0.05	0.06	0.00	0.00	0.03	0.08	0.33
PPE	8849	0.14	0.14	0.01	0.04	0.10	0.19	0.74

Table 1.2: CEO Tweeting Frequency around Earnings Announcements

Table 2 shows the OLS regression results of monthly number of tweets on the incidents of earnings announcements. The dependent variable is the log of one plus the monthly number of tweets made by a CEO. $Earnings_{t-1}$ is an indicator variable that equals one if the dependent variable is one month prior to the month of earnings announcement, and equals zero, otherwise. $Earnings_t$ and $Earnings_{t+1}$ are defined similarly. Firm controls are calculated one year prior to month t. CEO and year fixed effects are included. Standard errors are clustered by CEO and year. F-test is used to test if the coefficient of $Earnings_{t-1}$ is statistically different from each coefficient. F-stat values are shown below.

	(1)	(2)	(3)
		ln(1+Number of Tweets)	
$Earnings_{t-1}$	1.975***	0.939***	0.939***
	(32.91)	(17.82)	(17.82)
$Earnings_t$	1.920***	0.855***	0.855***
	(30.16)	(14.74)	(14.74)
$Earnings_{t+1}$	1.809***	0.795***	0.795***
	(28.56)	(14.53)	(14.53)
ROA			−0.309
			(−0.45)
MB			0.017
			(0.63)
Leverage			0.905**
			(2.98)
Ln(Assets)			−0.052***
			(−4.54)
Capex			0.174
			(0.68)
RD			−0.407**
			(−2.70)
PPE			−0.122
			(−1.21)
Constant	0.451***	5.648***	5.648***
	(13.54)	(5.40)	(5.40)
CEO Fixed Effects	NO	YES	YES
Year Fixed Effects	NO	YES	YES
R^2	0.333	0.676	0.686
Obs.	6,588	6,588	6,588
$H_0: Earnings_{t-1} = Earnings_t$			$F = 9.18$
$H_0: Earnings_{t-1} = Earnings_{t+1}$			$F = 27.46$

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.3: Pre-earnings Tweet Sentiment on Tweet Date CARs

Table 3 shows the regression results of CARs on daily tweet sentiment scores. CARs are calculated using the market model and Fama-French 3 Factor Model around the day on which a tweet is published. Panel A shows the regression results with the dependent variable as the CAR for [0,+1] where day 0 is the tweet date. Panel B shows the results with the dependent variable as the CAR for [-1,+1]. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values that ranges from -1 to 1. Tweet sentiment scores are aggregated daily. The sample includes tweets that are 30 days before the firm's earnings announcement. Tweet characteristics are logged to adjust for skewness. Firm characteristics are calculated one year prior to day t. Standard errors are clustered by CEO and year.

Panel A. Dependent Variable: CAR(0,+1) Around Tweet Day

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Tweet Day CAR(0,+1)				Tweet Day FF3(0,+1)			
Sentiment Score	0.190^{**}	0.223^{**}	0.222^{**}	0.252^{**}	0.170^{**}	0.215^{**}	0.215^{**}	0.239^{**}
	(2.33)	(2.19)	(2.19)	(2.41)	(2.16)	(2.35)	(2.35)	(2.53)
Ln(Likes)			−0.172	−0.165			−0.232	−0.219
			(−1.14)	(−1.09)			(−1.62)	(−1.51)
Ln(Replies)			0.950^{**}	0.906^{*}			0.909^{*}	0.859^{*}
			(2.05)	(1.87)			(1.88)	(1.74)
Ln(Retweets)			-0.066^{**}	-0.061^{**}			−0.040	−0.035
			(−2.24)	(−2.02)			(−1.33)	(−1.11)
ROA_{t-1}				−0.822				−0.882
				(−0.96)				(−1.06)
MB_{t-1}				-0.077^{*}				-0.086^{**}
				(−1.86)				(−2.21)
$Leverage_{t-1}$				0.085				0.180
				(0.22)				(0.53)
$Ln(Assets)_{t-1}$				−0.026				-0.028^{*}
				(−1.55)				(−1.76)
$Capex_{t-1}$				−0.436				−0.684
				(−0.62)				(−1.20)
RD_{t-1}				0.220				−0.030
				(0.33)				(−0.01)
PPE_{t-1}				−0.585				0.169
				(−0.24)				(0.07)
Constant	-0.080^{*}	-0.092^{**}	-0.102^{**}	2.515	−0.010	−0.027	−0.038	2.787^{*}
	(−1.75)	(−2.38)	(−2.60)	(1.58)	(−0.23)	(−0.79)	(−1.05)	(1.86)
CEO Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
Year Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
R^2	0.002	0.017	0.017	0.016	0.003	0.021	0.017	0.015
Obs.	8,849	8,844	8,844	8,263	8,849	8,844	8,844	8,263

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B. Dependent Variable: CAR(-1,+1) Around Tweet Day

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Tweet Day CAR(-1,+1)				Tweet Day FF3(-1,+1)			
Sentiment Score	0.266**	0.262**	0.261**	0.283**	0.244**	0.248**	0.248**	0.266**
	(2.47)	(2.10)	(2.10)	(2.26)	(2.44)	(2.28)	(2.28)	(2.43)
Ln(Likes)			−0.223	−0.193			−0.266*	−0.239*
			(−1.64)	(−1.52)			(−1.77)	(−1.69)
Ln(Replies)			0.141**	0.129*			0.132**	0.118*
			(2.16)	(1.94)			(1.99)	(1.79)
Ln(Retweets)			−0.096	−0.088			−0.078	−0.069
			(−1.55)	(−1.37)			(−1.23)	(−1.06)
Constant	−0.099	−0.097**	−0.114**	3.610	−0.005	−0.006	−0.021	3.224
	(−1.57)	(−2.05)	(−2.33)	(1.63)	(−0.08)	(−0.15)	(−0.49)	(1.54)
Firm Controls	NO	NO	NO	YES	NO	NO	NO	YES
CEO Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
Year Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
R^2	0.002	0.017	0.018	0.017	0.003	0.017	0.018	0.017
Obs.	8,849	8,844	8,844	8,263	8,849	8,844	8,844	8,263

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.4: Pre-earnings Negative and Positive Tweet Sentiment on Tweet Date CARs

Table 4 shows the regression results of CARs on daily tweet sentiment scores. CARs are calculated using the market model and Fama-French 3 Factor Model around the day on which a tweet is published. Panel A shows the regression results with the dependent variable as the CAR for [0,+1] where day 0 is the tweet date. Panel B shows the results with the dependent variable as the CAR for [-1,+1]. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values that ranges from -1 to 1. Sentiment scores are decomposed into Negative and Positive Scores, which range from 0 to 1, where 1 is the highest level. Tweet sentiment scores are aggregated daily. The sample includes tweets that are 30 days before the firm's earnings announcement. Tweet characteristics are logged to adjust for skewness. Firm characteristics are calculated one year prior to day t. Standard errors are clustered by CEO and year.

Panel A. Dependent Variable: CAR(0,+1) Around Tweet Day

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Tweet Day CAR(0,+1)				Tweet Day FF3(0,+1)			
Negative Score	-0.391^{**}	-0.283^{**}	-0.298^{**}	-0.324^{**}	-0.348^{***}	-0.301^{**}	-0.311^{**}	-0.358^{**}
	(-1.99)	(-2.01)	(-2.13)	(-2.13)	(-2.84)	(-2.02)	(-2.10)	(-2.14)
Positive Score	0.037^{**}	0.094^{***}	0.090^{***}	0.080^{***}	0.028	0.090^{***}	0.090^{***}	0.085^{***}
	(1.02)	(4.32)	(3.93)	(2.81)	(1.41)	(3.80)	(3.66)	(2.91)
Constant	0.000	0.000	0.000	0.031	0.001^{*}	0.001^{***}	0.001^{***}	0.037
	(0.33)	(0.23)	(0.01)	(1.34)	(1.78)	(3.86)	(3.74)	(1.55)
Tweet Controls	NO	NO	YES	YES	NO	NO	YES	YES
Firm Controls	NO	NO	NO	YES	NO	NO	NO	YES
CEO Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
Year Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
R^2	0.002	0.034	0.035	0.037	0.002	0.038	0.039	0.038
Obs.	8,849	8,844	8,844	8,263	8,849	8,844	8,844	8,263

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B. Dependent Variable: CAR(-1,+1) Around Tweet Day

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Tweet Day CAR(-1,+1)				Tweet Day FF3(-1,+1)			
Negative Score	-0.477^{***}	-0.286^{*}	-0.307^{**}	-0.289^{*}	-0.458^{***}	-0.318^{**}	-0.331^{**}	-0.343^{**}
	(-3.33)	(-1.95)	(-2.04)	(-1.75)	(-2.87)	(-2.37)	(-2.43)	(-2.24)
Positive Score	0.019	0.117^{**}	0.112^{**}	0.087^{*}	0.023	0.116^{***}	0.116^{***}	0.105^{***}
	(0.69)	(2.42)	(2.38)	(1.69)	(1.04)	(3.19)	(3.29)	(2.77)
Constant	0.000	-0.001^{***}	-0.001^{***}	0.043	0.001	0.000	0.000	0.024
	(0.48)	(-3.17)	(-2.90)	(1.25)	(1.20)	(0.21)	(0.73)	(0.67)
Tweet Controls	NO	NO	YES	YES	NO	NO	YES	YES
Firm Controls	NO	NO	NO	YES	NO	NO	NO	YES
CEO Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
Year Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
R^2	0.002	0.036	0.037	0.040	0.003	0.036	0.038	0.038
Obs.	8,849	8,844	8,844	8,263	8,849	8,844	8,844	8,263

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.5: Pre-earnings Tweet Sentiment and the Likelihood of Meeting Earnings Expectations

Table 5 shows the logit regression results of Beat on tweet sentiment scores. Beat is an indicator variable that equals one if a firm meets earnings expectations, and equals zero, otherwise. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values. Sentiment scores are decomposed into Negative and Positive Scores, which range from 0 to 1, where 1 is the highest level. Panel A uses the continuous variable, and Panel B uses an indicator variable that equals one if the scores are positive. For each firm's earnings announcement, tweet sentiment scores are aggregated on tweets made 30 days prior to the earnings announcement date. Firm controls are calculated one year prior to earnings announcement date. Standard errors are clustered by CEO and year.

Panel A. Continuous Sentiment Score

	(1)	(2)	(3)	(4)	(5)	(6)
			Beat			
Negative Score	1.175^{***}	1.225^{**}				
	(2.60)	(2.07)				
Positive Score			0.119	0.126		
			(0.56)	(0.28)		
Sentiment Score					-0.384^{***}	-0.470^{**}
					(−3.26)	(−2.14)
Log(Likes)		0.216		0.237		0.355
		(0.66)		(0.77)		(1.04)
Log(Replies)		0.012		0.064		0.036
		(0.08)		(0.49)		(0.18)
Log(Retweets)		0.299		0.308		0.251
		(1.17)		(1.16)		(0.95)
ROA		−0.128		-0.145^{*}		−0.163
		(−1.58)		(−1.73)		(−1.16)
MB		−0.046		−0.036		−0.029
		(−0.58)		(−0.44)		(−0.35)
Leverage		−0.348		−0.396		−0.022
		(−0.55)		(−0.60)		(−0.02)
Ln(Assets)		0.119^{**}		0.115^{**}		0.113^{***}
		(2.49)		(2.45)		(2.65)
Capex		0.355		0.347		0.005
		(0.34)		(0.33)		(0.01)
RD		1.757		1.820		1.489
		(0.72)		(0.72)		(0.24)
PPE		0.365		0.314		0.728
		(0.25)		(0.22)		(0.75)
Constant	0.949^{***}	-1.333^{*}	0.594^{***}	−1.307	1.349^{***}	−0.940
	(9.84)	(−1.80)	(5.86)	(−1.61)	(17.77)	(−1.21)
CEO Fixed Effects	NO	YES	NO	YES	NO	YES
Year Fixed Effects	NO	YES	NO	YES	NO	YES
Pseudo R^2	0.004	0.079	0.002	0.077	0.003	0.082
Obs.	1,137	1,137	1,137	1,137	1,137	1,137

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B. Indicator Sentiment Variable

	(1)	(2)	(3)	(4)	(5)	(6)
			Beat			
Negative	0.420***	0.435***			0.424***	0.444***
	(3.14)	(2.67)			(3.09)	(2.69)
Positive			0.148	0.022	−0.038	−0.117
			(0.58)	(0.07)	(−0.14)	(−0.36)
Constant	0.826***	−1.389	0.908***	−1.316	0.859***	−1.295
	(8.77)	(−0.97)	(3.67)	(−0.91)	(3.47)	(−0.89)
Firm Controls	NO	YES	NO	YES	NO	YES
CEO Fixed Effects	NO	YES	NO	YES	NO	YES
Year Fixed Effects	NO	YES	NO	YES	NO	YES
Pseudo R^2	0.005	0.082	0.004	0.076	0.007	0.082
Obs.	1,137	1,137	1,137	1,137	1,137	1,137

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.6: Tweet Sentiment and the Likelihood of Analyst Revising Forecasts Downwards

Table 6 shows the logit results of Downward Revision on tweet sentiment scores. Downward Revision is an indicator variable that equals one if an analyst revises his/her estimate downwards in the one month prior to earnings, and equals zero, otherwise. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values. Sentiment scores are decomposed into Negative and Positive Scores, which range from 0 to 1, where 1 is the highest level. For each firm's earnings announcement, tweet sentiment scores are aggregated on tweets made 60 days to 30 days prior to earnings announcement dates. Panel A shows the results using Downward Revision as the dependent variable. Panel B uses Upward Revision, which equals one if an analyst revises his/her estimate upwards in the one month prior to earnings. Panel C uses the number of downward revisions in the one month prior to earnings. Firm controls are calculated one year prior to earnings announcement date. Standard errors are clustered by CEO and year.

Panel A. Downward Revisions

	(1)	(2)	(3)	(4)	(5)	(6)
			Downward Revision			
Negative Score	2.463^{**}	2.773^{**}				
	(2.25)	(2.47)				
Positive Score			0.257	0.204		
			(0.96)	(0.78)		
Sentiment Score					-1.924^{**}	-2.165^{**}
					(−1.98)	(−2.07)
ROA		-0.303^{*}		-0.314^{*}		-0.314^{*}
		(−1.91)		(−1.92)		(−1.87)
MB		0.051		0.056		0.056^{**}
		(1.07)		(1.19)		(2.08)
Leverage		0.380		0.344		0.344
		(0.47)		(0.43)		(0.60)
Ln(Assets)		−0.106		−0.071		−0.071
		(−0.46)		(−0.31)		(−0.42)
Capex		−0.844		−0.803		−0.803
		(−1.49)		(−1.41)		(−1.36)
RD		−0.435		−0.311		−0.311
		(−0.11)		(−0.08)		(−0.07)
PPE		−1.112		−1.053		−1.053
		(−0.45)		(−0.42)		(−0.75)
Constant	0.486^{***}	0.275	0.465^{***}	0.211	0.452^{***}	0.211
	(30.57)	(0.84)	(25.82)	(0.58)	(23.82)	(0.64)
Tweet Controls	NO	YES	NO	YES	NO	YES
CEO Fixed Effects	NO	YES	NO	YES	NO	YES
Year Fixed Effects	NO	YES	NO	YES	NO	YES
Pseudo R^2	0.005	0.230	0.005	0.227	0.006	0.228
Obs.	1,617	1,617	1,617	1,617	1,617	1,617

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B. Upward Revisions

	(1)	(2)	(3)	(4)	(5)	(6)
			Upward Revision			
Negative Score	−1.067	−1.577				
	(−0.97)	(−1.42)				
Positive Score			−0.174	−0.104		
			(−0.70)	(−0.44)		
Sentiment Score					0.983	0.967
					(0.91)	(0.93)
Constant	−0.443	−0.390	−0.477	0.016	−0.512	0.016
	(−0.55)	(−0.11)	(−0.60)	(0.00)	(−0.64)	(0.00)
Tweet Controls	NO	YES	NO	YES	NO	YES
Firm Controls	NO	YES	NO	YES	NO	YES
CEO Fixed Effects	NO	YES	NO	YES	NO	YES
Year Fixed Effects	NO	YES	NO	YES	NO	YES
Pseudo R^2	0.003	0.140	0.003	0.136	0.004	0.137
Obs.	1,617	1,617	1,617	1,617	1,617	1,617

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel C. Number of Downward Revisions

	(1)	(2)	(3)	(4)	(5)	(6)
			Number of Downward Revision			
Negative Score	1.206*	1.421**				
	(2.12)	(2.48)				
Positive Score			0.194	0.181		
			(0.60)	(0.55)		
Sentiment Score					−1.103*	−1.207**
					(−1.90)	(−2.01)
Constant	0.174***	−0.539***	0.181***	−0.543***	0.158***	−0.543***
	(52.67)	(−15.93)	(15.64)	(−9.94)	(18.36)	(−15.62)
Tweet Controls	NO	YES	NO	YES	NO	YES
Firm Controls	NO	YES	NO	YES	NO	YES
CEO Fixed Effects	YES	YES	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES	YES	YES
R^2	0.022	0.328	0.019	0.330	0.067	0.331
Obs.	1,617	1,617	1,617	1,617	1,617	1,617

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.7: CEO Tweets Frequency around Earnings Announcements: Firm vs Non-firm Related

Table 7 shows the OLS regression results of monthly number of tweets on the incidents of earnings announcements. The dependent variable is the log of one plus the monthly number of tweets made by a CEO. $Earnings_{t-1}$ is an indicator variable that equals one if the dependent variable is one month prior to the month of earnings announcement, and equals zero, otherwise. $Earnings_t$ and $Earnings_{t+1}$ are defined similarly. Panel A shows the subsample with firm-related tweets. Panel B shows the subsample with non-firm related tweets. Firm controls are calculated one year prior to month t. CEO and year fixed effects are included. Standard errors are clustered by CEO and year. F-test is used to test if the coefficient of $Earnings_{t-1}$ is statistically different from each coefficient. F-stat values are shown below.

Panel A. Firm-Related Tweets

	(1)	(2)	(3)	(4)
	ln(1+Number of Tweets)			
$Earnings_{t-1}$	1.010***	0.605***	0.594***	0.592***
	(37.40)	(21.80)	(21.91)	(21.80)
$Earnings_t$	0.708***	0.318***	0.313***	0.306***
	(20.27)	(10.90)	(11.01)	(10.69)
$Earnings_{t+1}$	0.706***	0.299***	0.291***	0.285***
	(20.83)	(10.04)	(10.02)	(9.84)
Constant	0.330***	0.509***	1.377**	1.500**
	(18.15)	(33.95)	(2.74)	(3.06)
CEO Fixed Effects	NO	YES	YES	YES
Year Fixed Effects	NO	YES	YES	YES
R^2	0.919	0.920	0.920	0.924
Obs.	6,588	6,588	6,588	6,588
$H_0: Earnings_{t-1} = Earnings_t$	$F = 103.36$			
$H_0: Earnings_{t-1} = Earnings_{t+1}$	$F = 131.40$			

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B. Non Firm-Related Tweets

	(1)	(2)	(3)	(4)
	ln(1+Number of Tweets)			
$Earnings_{t-1}$	1.439***	0.673***	0.644***	0.642***
	(29.99)	(16.62)	(16.95)	(16.86)
$Earnings_t$	1.137***	0.426***	0.399***	0.397***
	(22.21)	(11.43)	(11.49)	(11.44)
$Earnings_{t+1}$	1.181***	0.424***	0.397***	0.394***
	(21.87)	(10.77)	(10.65)	(10.56)
Constant	−0.195***	−0.314***	0.0764	0.139
	(−8.21)	(−7.03)	(0.07)	(0.13)
CEO Fixed Effects	NO	YES	YES	YES
Year Fixed Effects	NO	YES	YES	YES
R^2	0.835	0.840	0.843	0.844
Obs.	6,588	6,588	6,588	6,588
$H_0: Earnings_{t-1} = Earnings_t$	$F = 66.09$			
$H_0: Earnings_{t-1} = Earnings_{t+1}$	$F = 66.24$			

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.8: Firm-Related vs Non-Firm Related Pre-earnings Tweet Sentiment on Tweet Date CARs

Table 8 shows the regression results of CARs on daily tweet sentiment scores. CARs are calculated using the market model and Fama-French 3 Factor Model around the day in which a tweet is published. Panel A shows the regression results with the dependent variable as the CAR for [0,+1] where day 0 is the tweet date. Panel B shows the results with the dependent variable as the CAR for [-1,+1]. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values that ranges from -1 to 1. Tweet sentiment scores are aggregated daily. The sample includes tweets that are 30 days before the firm's earnings announcement. The sample is split into sub-samples, containing firm-related tweets and non-firm related tweets, using stochastic gradient descent. Tweet characteristics are logged to adjust for skewness. Firm characteristics are calculated one year prior to day t. Standard errors are clustered by CEO and year.

Panel A. Dependent Variable: CAR(0,+1)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Tweet Day CAR(0,+1)							
	Firm-Related				Non-firm Related			
Sentiment Score	0.292***	0.311**	0.312**	0.340**	0.249**	0.236**	0.234**	0.250**
	(2.70)	(2.40)	(2.42)	(2.56)	(2.57)	(2.13)	(2.11)	(2.18)
Constant	−0.053	−0.059	−0.118*	3.245	−0.088*	−0.084	−0.082	3.925*
	(−0.90)	(−1.02)	(−1.85)	(1.33)	(−1.70)	(−1.61)	(−1.54)	(1.96)
Tweet Controls	NO	NO	YES	YES	NO	NO	YES	YES
Firm Controls	NO	NO	NO	YES	NO	NO	NO	YES
CEO Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
R^2	0.005	0.020	0.024	0.024	0.004	0.021	0.021	0.020
Obs.	4,908	4,908	4,908	4,908	6,938	6,938	6,938	6,938

Panel B. Dependent Variable: CAR(-1,+1)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Tweet Day CAR(0,+1)							
	Firm-Related				Non-firm Related			
Sentiment Score	0.394***	0.377**	0.379**	0.384**	0.353***	0.299**	0.297**	0.295**
	(3.04)	(2.52)	(2.53)	(2.49)	(3.09)	(2.31)	(2.30)	(2.22)
Constant	−0.091	−0.085	−0.158**	3.264	−0.113*	−0.094	−0.095	5.154*
	(−1.28)	(−1.24)	(−2.14)	(1.03)	(−1.73)	(−1.48)	(−1.46)	(1.91)
Tweet Controls	NO	NO	YES	YES	NO	NO	YES	YES
Firm Controls	NO	NO	NO	YES	NO	NO	NO	YES
CEO Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
R^2	0.005	0.019	0.025	0.025	0.003	0.019	0.020	0.020
Obs.	4,908	4,908	4,908	4,908	6,938	6,938	6,938	6,938

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.9: Pre-earnings Tweet Sentiment and the Likelihood of Meeting Earnings Expectations: Firm vs Non-firm Related

Table 9 shows the logit regression results of Beat on tweet sentiment scores. Beat is an indicator variable that equals one if a firm meets earnings expectations, and equals zero, otherwise. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values. Sentiment scores are decomposed into Negative and Positive Scores, which range from 0 to 1, where 1 is the highest level. Panel A uses the continuous variable, and Panel B uses an indicator variable that equals one if the scores are positive. The sample is split into sub-samples, containing firm-related tweets and non-firm related tweets, using stochastic gradient descent. For each firm's earnings announcement, tweet sentiment scores are aggregated on tweets made 30 days prior to the earnings announcement date. Firm controls are calculated one year prior to earnings announcement date. Standard errors are clustered by CEO and year.

	(1)	(2)	(3)	(4)	(5)	(6)
			Beat			
		Firm-Related			Non-firm Related	
Negative Score	1.551***			1.531***		
	(3.07)			(2.88)		
Positive Score		0.247			0.772	
		(0.54)			(1.64)	
Sentiment Score			−0.538**			−0.252
			(−2.24)			(−1.03)
Constant	−1.101	−1.100	−0.697	−0.382**	−0.369**	−0.674***
	(−1.33)	(−1.27)	(−0.96)	(−2.28)	(−2.26)	(−2.71)
Firm Controls	NO	YES	NO	YES	NO	YES
Pseudo R^2	0.026	0.022	0.029	0.087	0.026	0.022
Obs.	1,112	1,112	1,112	1,112	1,112	1,112

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.10: Tweet Sentiment and the Likelihood of Analyst Revising Forecasts Downwards: Firm vs Non-firm Related

Table 10 shows the logit results of Downward Revision on tweet sentiment scores. Downward Revision is an indicator variable that equals one if an analyst revises his/her estimate downwards in the one month prior to earnings, and equals zero, otherwise. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values. Sentiment scores are decomposed into Negative and Positive Scores, which range from 0 to 1, where 1 is the highest level. For each firm's earnings announcement, tweet sentiment scores are aggregated on tweets made 60 days to 30 days prior to earnings announcement dates. Panel A shows the results using Downward Revision as the dependent variable. Panel B uses Upward Revision, which equals one if an analyst revises his/her estimate upwards in the one month prior to earnings. Panel C uses the number of downward revisions in the one month prior to earnings. The sample is split into sub-samples, containing firm-related tweets and non-firm related tweets, using stochastic gradient descent. Firm controls are calculated one year prior to earnings announcement date. Standard errors are clustered by CEO and year.

	(1)	(2)	(3)	(4)	(5)	(6)
	Downward Revision					
	Firm-Related			Non-firm Related		
Negative Score	2.923**			1.889*		
	(2.18)			(1.88)		
Positive Score		0.029			0.099	
		(0.05)			(0.38)	
Sentiment Score			−3.067*			−1.33
			(−1.90)			(−1.29)
Constant	2.405	2.425	2.424	0.717	0.535	0.535
	(0.70)	(0.68)	(0.81)	(0.26)	(0.20)	(0.34)
Tweet Controls	NO	YES	NO	YES	NO	YES
Firm Controls	NO	YES	NO	YES	NO	YES
CEO Fixed Effects	YES	YES	YES	YES	YES	YES
Pseudo R^2	0.251	0.250	0.250	0.230	0.228	0.227
Obs.	1,617	1,617	1,617	1,617	1,617	1,617

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.11: Tweet Sentiment and Account Age

Table 11 shows the logit regression results of Beat on tweet sentiment scores interacted with account age. Beat is an indicator variable that equals one if a firm meets earnings expectations, and equals zero, otherwise. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values. Sentiment scores are decomposed into Negative and Positive Scores, which range from 0 to 1, where 1 is the highest level. For each firm's earnings announcement, tweet sentiment scores are aggregated on tweets made 30 days prior to the earnings announcement date. Account Age is the number of months at the date of earnings announcement from the first tweet made in the CEO's account. Firm controls are calculated one year prior to earnings announcement date. Standard errors are clustered by CEO and year.

Continuous Sentiment Score

	(1)	(2)	(3)	(4)	(5)	(6)
			Beat			
Negative Score × Account Age	0.086	0.060				
	(1.35)	(0.86)				
Negative Score	0.846**	0.916*				
	(2.02)	(1.96)				
Positive Score × Account Age			0.013*	0.009		
			(1.95)	(0.64)		
Positive Score			0.156	0.370		
			(1.41)	(0.55)		
Sentiment Score × Account Age					−0.025	−0.018
					(−1.06)	(−0.58)
Sentiment Score					−0.317*	−0.603**
					(−1.87)	(−2.16)
Account Age	0.022***	0.032***	0.011***	0.030***	0.021***	0.031***
	(4.08)	(2.97)	(2.93)	(2.98)	(2.91)	(3.10)
Log(Likes)		0.168		0.098		0.104
		(0.38)		(0.24)		(0.25)
Log(Replies)		−0.239		−0.364		−0.356
		(−1.09)		(−1.48)		(−1.46)
Log(Retweets)		0.187		0.123		0.127
		(0.78)		(0.46)		(0.47)
ROA		−0.511		−0.137		−0.155
		(−0.68)		(−0.18)		(−0.20)
MB		−0.061		−0.062		−0.059
		(−0.69)		(−0.70)		(−0.67)
Leverage		0.082		0.113		0.098
		(0.13)		(0.18)		(0.15)
Ln(Assets)		0.158***		0.153***		0.152***
		(3.59)		(3.54)		(3.47)
Capex		0.197		0.284		0.284
		(0.22)		(0.31)		(0.31)
RD		0.293		0.332		0.335
		(1.19)		(1.55)		(1.53)
PPE		−0.147		−0.095		−0.113
		(−0.10)		(−0.07)		(−0.08)
Constant	−0.481	−1.839***	0.714***	−1.934***	0.721***	−1.932***
	(−0.96)	(−3.27)	(4.02)	(−3.11)	(4.10)	(−3.10)
CEO Fixed Effects	NO	YES	NO	YES	NO	YES
Year Fixed Effects	NO	YES	NO	YES	NO	YES
Pseudo R^2	0.004	0.079	0.002	0.077	0.003	0.082
Obs.	1,137	1,137	1,137	1,137	1,137	1,137

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.12: Extreme Pre-earnings Tweet Sentiment on Tweet Date CARs

Table 12 shows the regression results of CARs on daily extreme tweet sentiment scores. CARs are calculated using the market model and Fama-French 3 Factor Model around the day on which a tweet is published. The dependent variable is the CAR for [0,+1] where day 0 is the tweet date. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values that ranges from -1 to 1. Only the sentiment score from the maximum negative or positive tweet are used per tweet day. Tweet sentiment scores are aggregated daily. The sample includes tweets that are 30 days before the firm's earnings announcement. Tweet characteristics are logged to adjust for skewness. Firm characteristics are calculated one year prior to day t. Standard errors are clustered by CEO and year.

Dependent Variable: CAR(0,+1) Around Tweet Day

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Tweet Day CAR(0,+1)				Tweet Day FF3(0,+1)			
Max Negative	−0.542	−0.518	−0.548	−0.503	−0.501	−0.553	−0.578	−0.576
	(−1.15)	(−0.80)	(−0.82)	(−0.71)	(−1.07)	(−0.84)	(−0.85)	(−0.78)
Max Positive	0.246	0.368	0.347	0.421	0.160	0.282	0.264	0.343
	(1.22)	(1.35)	(1.25)	(1.35)	(0.78)	(1.07)	(0.99)	(1.14)
Ln(Likes)			−0.385	−0.348			−0.397	−0.358
			(−1.03)	(−0.86)			(−1.27)	(−1.02)
Log(Replies)			11.342	11.177			9.400	9.197
			(1.55)	(1.47)			(1.35)	(1.31)
Ln(Retweets)			−0.571	−0.551			−0.277	−0.258
			(−1.12)	(−1.00)			(−0.56)	(−0.49)
ROA_{t-1}				−0.502				−0.305
				(−0.36)				(−0.20)
MB_{t-1}				-0.129^{*}				-0.143^{*}
				(−1.95)				(−1.83)
$Leverage_{t-1}$				0.228				0.454
				(0.41)				(0.83)
$Ln(Assets)_{t-1}$				-0.630^{**}				-0.632^{**}
				(−2.12)				(−2.06)
$Capex_{t-1}$				−1.165				-1.264^{*}
				(−1.27)				(−1.69)
RD_{t-1}				1.683				−0.627
				(0.19)				(−0.07)
PPE_{t-1}				−1.827				−1.395
				(−0.50)				(−0.36)
Constant	−0.034	−0.069	−0.071	6.127^{**}	0.051	0.021	0.019	6.298^{**}
	(−0.48)	(−0.79)	(−0.86)	(2.12)	(0.78)	(0.24)	(0.23)	(2.09)
CEO Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
Year Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
R^2	0.002	0.025	0.026	0.027	0.003	0.025	0.025	0.026
Obs.	8,849	8,844	8,844	8,263	8,849	8,844	8,844	8,263

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.13: Pre-earnings Tweet Sentiment on Tweet Date CARs: Single Tweet Days

Table 13 shows the regression results of CARs on daily tweet sentiment scores. CARs are calculated using the market model and Fama-French 3 Factor Model around the day on which a tweet is published. The dependent variable is the CAR for [0,+1] where day 0 is the tweet date. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values that ranges from -1 to 1. Only tweet days with exactly one tweet made are included in this sample. The sample includes tweets that are 30 days before the firm's earnings announcement. Tweet characteristics are logged to adjust for skewness. Firm characteristics are calculated one year prior to day t. Standard errors are clustered by CEO and year.

Dependent Variable: CAR(0,+1) Around Tweet Day

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Tweet Day CAR(0,+1)				Tweet Day FF3(0,+1)			
Negative	-0.440^{**}	-0.537^{**}	-0.548^{**}	-0.468^{**}	-0.454^{**}	-0.551^{***}	-0.561^{***}	-0.492^{***}
	(−2.08)	(−2.50)	(−2.53)	(−2.51)	(−2.20)	(−2.63)	(−2.67)	(−2.64)
Positive	0.011	−0.016	−0.021	0.012	−0.007	−0.034	−0.039	−0.008
	(0.12)	(−0.16)	(−0.21)	(0.12)	(−0.07)	(−0.34)	(−0.39)	(−0.08)
Log(Likes)			-0.787^{***}	-0.696^{*}			-0.776^{***}	-0.582^{*}
			(−3.18)	(−1.73)			(−3.37)	(−1.87)
Log(Replies)			5.284	6.761			2.113	3.208
			(0.79)	(1.04)			(0.28)	(0.44)
Log(Retweets)			−0.371	−0.583			0.113	−0.118
			(−0.59)	(−0.88)			(0.18)	(−0.18)
ROA_{t-1}				−0.357				−0.362
				(−0.90)				(−0.88)
MB_{t-1}				−0.104				−0.081
				(−0.85)				(−0.70)
$Leverage_{t-1}$				−0.049				−0.731
				(−0.03)				(−0.43)
$Ln(Assets)_{t-1}$				-0.155^{***}				-0.153^{**}
				(−2.68)				(−2.49)
$Capex_{t-1}$				−0.791				−0.458
				(−0.51)				(−0.30)
RD_{t-1}				0.116				0.070
				(1.42)				(0.87)
PPE_{t-1}				−0.654				−0.467
				(−0.76)				(−0.56)
Constant	−0.060	0.021	0.052	14.367^{**}	0.062	0.146	0.179	14.356^{**}
	(−0.24)	(0.09)	(0.22)	(2.35)	(0.25)	(0.61)	(0.75)	(2.27)
CEO Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
Year Fixed Effects	NO	YES	YES	YES	NO	YES	YES	YES
R^2	0.004	0.095	0.096	0.100	0.004	0.095	0.096	0.096
Obs.	1,381	1,368	1,368	1,292	1,381	1,368	1,368	1,292

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.14: Falsification Test: Pre-earnings Tweet Sentiment and the Likelihood of Meeting Earnings Expectations

Table 14 shows the logit regression results of Beat Prior Quarter on tweet sentiment scores. Beat Prior Quarter is an indicator variable that equals one if a firm meets earnings expectations from the prior quarter, and equals zero, otherwise. Earnings Expectations is calculated using the prior quarter's earnings per share as the benchmark. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values. Sentiment scores are decomposed into Negative and Positive Scores, which range from 0 to 1, where 1 is the highest level. For each firm's earnings announcement, tweet sentiment scores are aggregated on tweets made 30 days prior to the earnings announcement date. Firm controls are calculated one year prior to earnings announcement date. Standard errors are clustered by CEO and year.

Continuous Sentiment Score

	(1)	(2)	(3)	(4)	(5)	(6)
			Beat Prior Quarter			
Negative Score	−0.097 (−0.72)	0.267 (0.95)				
Positive Score			−0.033 (−1.40)	0.151 (1.34)		
Sentiment Score					−0.026 (−0.74)	0.115 (1.09)
Log(Likes)		−0.176 (−0.55)		−0.193 (−0.62)		−0.188 (−0.62)
Log(Replies)		−0.167 (−0.82)		−0.200 (−1.05)		−0.197 (−1.03)
Log(Retweets)		0.294 (0.92)		0.268 (0.85)		0.267 (0.85)
ROA		−0.245 (−1.58)		−0.235 (−1.48)		−0.233 (−1.47)
MB		0.104* (1.90)		0.109* (1.84)		0.108* (1.84)
Leverage		0.342 (0.42)		0.314 (0.39)		0.298 (0.36)
Ln(Assets)		−0.185 (−0.64)		−0.186 (−0.62)		−0.197 (−0.65)
Capex		−1.207*** (−2.80)		−1.232*** (−2.84)		−1.247*** (−2.89)
RD		−0.706* (−1.72)		−0.681 (−1.59)		−0.702 (−1.62)
PPE		0.499 (0.25)		0.534 (0.26)		0.499 (0.25)
Constant	0.949*** (9.84)	−1.333* (−1.80)	0.594*** (5.86)	−1.307 (−1.61)	1.349*** (17.77)	−0.940 (−1.21)
CEO Fixed Effects	NO	YES	NO	YES	NO	YES
Year Fixed Effects	NO	YES	NO	YES	NO	YES
Pseudo R^2	0.004	0.079	0.002	0.077	0.003	0.082
Obs.	1,137	1,137	1,137	1,137	1,137	1,137

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.15: Pre-earnings Tweet Sentiment and the Likelihood of Meeting Earnings Expectations Marginally

Table 15 shows the logit regression results of Close Beat on tweet sentiment scores. Close Beat is an indicator variable that equals one if a firm meets earnings expectations marginally, and equals zero, otherwise. Marginally is defined if the actual EPS minus the mean analyst estimate in the prior month, divided by log price equals 0.01. Sentiment word scores are computed with Valence Aware Dictionary and sEntiment Reasoner (VADER), which contains a list of lexicon and their corresponding sentiment values. Sentiment scores are decomposed into Negative and Positive Scores, which range from 0 to 1, where 1 is the highest level. For each firm's earnings announcement, tweet sentiment scores are aggregated on tweets made 30 days prior to the earnings announcement date. Firm controls are calculated one year prior to earnings announcement date. Standard errors are clustered by CEO and year.

Continuous Sentiment Score

	(1)	(2)	(3)	(4)	(5)	(6)
			Close Beat			
Negative Score	1.132**	1.018**				
	(2.51)	(2.13)				
Positive Score			0.187	0.234		
			(1.60)	(1.57)		
Sentiment Score					−0.488*	−0.479**
					(−1.95)	(−2.17)
Log(Likes)		0.269		0.224		0.229
		(0.66)		(0.71)		(0.60)
Log(Replies)		−0.102		−0.190		−0.180
		(−0.56)		(−0.91)		(−0.96)
Log(Retweets)		0.272		0.223		0.225
		(1.18)		(0.88)		(0.87)
ROA		−0.134		−0.106		−0.109
		(−1.06)		(−0.58)		(−0.92)
MB		−0.013		−0.014		−0.012
		(−0.16)		(−0.14)		(−0.14)
Leverage		−0.409		−0.368		−0.380
		(−0.74)		(−0.38)		(−0.66)
Ln(Assets)		0.104***		0.097		0.097***
		(2.99)		(1.21)		(3.03)
Capex		0.284		0.349		0.347
		(0.37)		(0.40)		(0.44)
RD		0.707		1.042		1.039
		(0.21)		(0.31)		(0.32)
PPE		−0.632		−0.617		−0.624
		(−0.37)		(−0.52)		(−0.37)
Constant	0.791***	−1.399**	0.684***	−1.448	−0.701*	−1.444**
	(8.30)	(−2.23)	(4.41)	(−1.40)	(−1.71)	(−2.24)
CEO Fixed Effects	NO	YES	NO	YES	NO	YES
Year Fixed Effects	NO	YES	NO	YES	NO	YES
Pseudo R^2	0.004	0.079	0.002	0.077	0.003	0.082
Obs.	1,137	1,137	1,137	1,137	1,137	1,137

t-statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Chapter 2

Generalized Bayesian Inference on the Expectile Risk Measure

2.1 Introduction

Risk management is a core competence of financial institutions and techniques for quantifying market risk are central to the process of managing risk. Value-at-risk (VaR), the most popular measure to evaluate the market risk of a portfolio, identifies the loss that is likely to be exceeded by a specified probability such as 0.99. It has become the standard measure of market risk, and has been used by financial institutions over the past two decades for setting regulatory capital requirements. However, it is generally agreed that VaR has some drawbacks. For instance, VaR has been criticized for disregarding the magnitude of the loss beyond a designated threshold. In addition, VaR also lacks subadditivity and thus fails to satisfy mathematical principles characterizing coherent risk measures, as demonstrated by Artzner et al. (1999). Expected shortfall (ES), also called conditional VaR, is a risk measure that overcomes these weaknesses. ES is defined as the conditional expectation of the loss given that it exceeds the VaR and thus provides information on the magnitude of the loss beyond the VaR. It becomes increasingly widely used in practice. Nevertheless, ES fails the elicitability criterion that is a key property for a risk measure as it provides a natural methodology to perform backtesting. See Bellini and Bignozzi (2015), Ziegel (2016), and Chen (2018) for references.

Expectile, introduced by Newey and Powell (1987) in the context of linear regression, has been revealed as a reasonable risk measure to offset the weaknesses of VaR and ES. As pointed by Ziegel (2016), the expectile is a coherent and elicitable risk measure and it is the only risk measure satisfying these two properties. Therefore, the expectile is becoming an increasingly popular tool in risk management and capital allocation for financial institutions. In addition, the expectile has also been widely used as a tool for efficient estimation of the ES through a one-one mapping between the two, see Taylor (2008), Xie et al (2014) and Daouia et al (2018). It is well known that estimation of the expectile becomes straightforward when the population distribution of the market loss is parametric. For instance, Krätschmer and Zähle (2017) consider estimating the population expectile when the loss follows the log-normal distribution or the Pareto distribution. See other related work by Holzmann and Klar (2016) and Daouia et al. (2018), to mention just a few. Nevertheless, the parametric estimation of the expectile suffers from the issue of possible model misspecification and thus could result in severe estimation bias, as discussed by many researchers such as Uppal and Wang (2003), Schmeiser et al. (2012) and Hong and Martin (2020). The main purpose of this paper is to propose an approach for making inference on the population expectile that could avoid this model misspecification issue.

The contribution of this paper has three-fold: first, we build a detailed framework for making inference on the population expectile of the financial loss through the generalized Bayesian approach developed by Bissiri et al. (2016), which provide a full picture of the population expectile; second, we provide some theoretic support for our methodology; third, our approach avoids the potential model misspecification issue and thus might be applied to estimating other financial risk measures.

The rest of the paper is organized as follows. In Section 2 we briefly review the literature on the expectile and the generalized Bayesian approach. Section 3 introduces our methodology in detail, along with some theoretical support. A simulation study is carried out to validate the proposed methodology in Section 4. We also analyze one financial loss data to illustrate our new approach. Section 5 provides a summary and concluding comments.

2.2 Literature Review

2.2.1 Expectiles and their properties

Let Y denote a random variable for the financial loss of a portfolio. Assume that the positive values of Y stand for financial losses and the negative values for the gains. If one is only interested in the financial losses, we may simply assume that Y takes nonnegative values. For simplicity, we assume through this paper that Y is continuous with the cumulative distribution function $F(y)$ and the probability density function $f(y)$ unless specified otherwise. For any $\tau \in (0,1)$, the τ-expectile of Y can be uniquely be defined by

$$\begin{aligned} \mu_\tau &= \underset{t \in R}{\operatorname{argmin}} \left\{ \tau E[(Y-t)_+^2] + (1-\tau) E[(t-Y)_+^2] \right\} \\ &= \underset{t \in R}{\operatorname{argmin}} \ E[D_\tau(Y-t)] \end{aligned} \tag{2.1}$$

where $x_+ = \max\{x, 0\}$ and

$$D_\tau(x) = \begin{cases} \tau x^2, & x \geq 0, \\ (1-\tau)x^2, & x < 0, \end{cases} \tag{2.2}$$

when assuming that the second moment of Y exists. Clearly $\mu_{1/2}$ is just the usual mean of Y and thus the τ-expectile is a generalization of the mean of a random variable. The above definition of the expectile was first introduced by Newey and Powell (1987) in the context of linear regression. Note that the well-known τ-quantile, also called VaR, can be obtained similarly by replacing the asymmetrically weighted squared deviations in (2.2) by the asymmetrically weighted mean absolute deviations. It is noted also that the ES is defined as the conditional expectation of the loss given that it exceeds the VaR (see Yamai et al. (2002)).

One good property of the expectile can be observed from the first order conditions on the optimization in (2.1), namely, μ_τ is the unique solution of

$$\frac{\tau}{1-\tau} = \frac{E[(\mu_\tau - Y)\mathbf{1}(Y \leq \mu_\tau)]}{E[(Y-\mu_\tau)\mathbf{1}(Y \geq \mu_\tau)]}. \tag{2.3}$$

when assuming only that the first moment of Y exists. See Newey and Powell (1987) and Bellini and Di Bernardino (2017) and references therein for further properties of the expectile risk measure.

In particular, while VaR and ES focus only on extreme losses, the expectile balances gains and losses, which is desirable for portfolio management. Because of this, the expectile is also closely related to the Omega performance measure of Keating and Shadwick (2002). Bellini and Di Bernardino (2017) point out that the expectile has an intuitive interpretation in terms of its acceptance set, namely, a risk is acceptable (for a regulator) if its gains-loss ratio is sufficiently high. It is worth noting that expectiles generalize the expectation just as quantiles generalize the median and the expectile function summarizes the loss distribution in much the same way as the quantile function. The expectile represents a linguistic portmanteau of expectation and quantile. For other interpretations, see Martin (2014), Chen (2018) and Philipps (2021) in detail.

2.2.2 Generalized Bayesian Inference

Bayesian approach is well suited for making inference on the financial risk measures as it allows for a consistent and convenient statistical framework to quantify the uncertainties involved. Suppose that we have a random sample $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ from the population with the probability density function $f(x \mid \theta)$, where θ takes values in a parameter set Θ. If $\pi(\theta)$ represents prior beliefs about the unknown parameter θ, then the well-known Bayes' theorem tells us that the posterior distribution of θ can be obtained

$$\pi(\theta \mid \$\mathbf{x}) \;=\; \frac{\pi(\theta)\prod_{i=1}^{n} f(x_i \mid \theta)}{\int_{\Theta} \pi(\theta)\prod_{i=1}^{n} f(x_i \mid \theta)\, d\theta},$$

which can be viewed as an update of the prior belief for the unknown parameter θ through using the information in the observed data $\mathbf{x}$. The posterior distribution of the unknown quantity of interest, such as the τ-expectile of X, can be obtained straightforward. It is noted that this posterior distribution provides a full picture for the unknown quantity of the population using the information both in the data and in the prior belief. Bayesian statistics has now gained greater acceptance in financial modelling, see Rachev et al. (2008) and Jacquier and Polson (2011). Nevertheless, the above traditional Bayesian analysis heavily depends on the population distribution $f(x \mid \theta)$ specified. If $f(x \mid \theta)$ is misspecified, the posterior inference drawn can be quite misleading, see Uppal and Wang (2003), Schmesier et al. (2012) and Hong and Martin (2020).

To avoid the misspecification issue of the population distribution, Bissiri et al. (2016) propose a framework for generalized Bayesian inference on the unknown quantity of interest directly. For convenience, let's assume that the parameter θ itself is the unknown quantity of interest. Suppose there is a loss function $\ell_\theta(x)$ that links data to the parameter such that interest is in the θ minimizing $E[\ell_\theta(X)]$, where $E[\ell_\theta(X)]$ is often called the discrepancy or risk function in statistical literature. Bissiri et al. (2016) argue that a valid and coherent update of the prior belief $\pi(\theta)$ is given by

$$\pi(\theta \mid \mathbf{X}) \;=\; \frac{\pi(\theta)\exp\{-\omega n D_n(\theta)\}}{\int_{\Theta} \pi(\theta)\exp\{-\omega n D_n(\theta)\}\, d\theta}, \tag{2.4}$$

where ω is a tuning (also called scale) parameter to be determined and $D_n(\theta) = \sum_{i=1}^{n} \ell_\theta(x_i)$/n is the empirical version of the discrepancy function

$E[\ell_\theta(X)]$. We name (**??**) as the generalized posterior distribution of θ, while some researchers also call as the Gibbs posterior, see Syring et al. (2019). Clearly, the so-called generalized posterior coincides with the usual posterior distribution stated in the above when the tuning parameter $\omega = 1$ and the loss function is the negative average log-likelihood. Note also that the generalized Bayesian inference does not need to specify the population distribution $f(x \mid \theta)$ and thus is quite convenient for making inference on financial risk measures.

2.3 Generalized Posterior Distributions

2.3.1 Construction of the Posterior

Let $\mathbf{y} = (y_1, y_2, \ldots, y_n)$ be independent and identically distributed losses generated from an unknown loss distribution $F(y)$ with τ-expectile being denoted as μ_τ, defined by (2.2). The empirical discrepancy function is

$$D_n(\mu_\tau) \;=\; \frac{1}{n}\sum_{i=1}^{n} D_\tau(y_i - \mu_\tau), \tag{2.5}$$

where $D_\tau(x)$ is given by (2.2).

Suppose that we have a prior distribution $\pi(\mu_\tau)$ on the τ-expectile μ_τ. Combining the empirical discrepancy function $D_n(\mu_\tau)$ with a prior distribution $\pi(\mu_\tau)$, we can then obtain a generalized posterior density for μ_τ as follows

$$\pi(\mu_\tau \mid \mathbf{y}) \;=\; \frac{\pi(\mu_\tau)\exp\{-\omega n D_n(\mu_\tau)\}}{\int_\Theta \pi(\mu_\tau)\exp\{-\omega n D_n(\mu_\tau)\}\, d\mu_\tau}, \tag{2.6}$$

where Θ is the support of the τ-expectile μ_τ and ω is a tuning parameter to be determined. Note that the empirical discrepancy function $D_n(\mu_\tau)$ increases quadratically for large μ_τ and then the negative sign in the exponent implies that the integrant in (2.6) is indeed integrable for any reasonable prior $\pi(\mu_\tau)$. Therefore, the generalized posterior $\pi(\mu_\tau \mid \mathbf{y})$ is well-defined on the support Θ. Since it is only one-dimensional, computation of any relevant feature of the generalized posterior is straightforward.

As for the prior distribution, since the τ-expectile is a practically meaningful quantity, the researcher may have genuine prior information available (e.g. based on historical data) from which an informative prior distribution can be constructed. Note that the prior distribution is simply a prior belief on the τ-expectile and thus it is not required to be very accurate. As a general recommendation, we suggest a gamma prior distribution with shape and scale parameters chosen so that the mean reflects some genuine prior information and the standard deviation is some (potentially large) fraction of that mean; see next section for more details on this prior specification when dealing with the analysis of the real data. Note that also $\pi(\mu_\tau \mid \mathbf{y})$ has a tuning parameter ω involved. Bissiri et al. (2016) provided several suggestions for the choice of ω, including using annealing, unit information loss, hierachical loss and operational characteristics and subjective calibration. We recommend choosing the tuning parameter using the unit information loss due to its simplicity and now describe this approach below.

Let $\hat{\mu}_\tau^*$ be the prior mode. Using the unit information loss matching between the prior and the data, Bissiri et al. (2016) explained that the tuning parameter ω can be determined by

$$\omega \;=\; \frac{E_{\mu_\tau}[\log\{\pi(\hat{\mu}_\tau^*)/\pi(\mu_\tau)\}]}{E_{\mu_\tau,Y}[D_\tau(Y-\mu_\tau)]},$$

which immediately suggests that the tuning parameter can be estimated by

$$\hat{\omega} \;=\; \frac{\int \pi(\mu_\tau)\log\frac{\pi(\hat{\mu}_\tau^*)}{\pi(\mu_\tau)}\;d\mu_\tau}{\frac{1}{n}\sum_{i=1}^n D_\tau(y_i-\hat{\mu}_\tau)},$$

where $\hat{\mu}_\tau$ is the empirical estimate of μ_τ that minimizes the discrepancy function (2.5).

2.3.2 Posterior Consistency

We now investigate the asymptotic behaviour of the generalized posterior distribution $\pi(\mu_\tau \mid \mathbf{y})$ in (2.6) of the τ-expectile μ_τ, namely, consistency, when the tuning parameter ω is assumed fixed. Note that consistency is a desirable property for any statistical procedure. Let $\mu_\tau^* = \mu_\tau^*(P)$ denote the true τ-expectile of the population distribution P. We say that the generalized posterior distribution $\pi(\mu_\tau \mid \mathbf{y})$ is consistent if

$$Pr(\{\mu_\tau : |\mu_\tau - \mu_\tau^*(P)| \le \epsilon\} \mid y_1, y_2, \ldots, y_n) \to 1 \text{ as } n \to \infty \qquad (2.7)$$

or equivalently

$$Pr(\{\mu_\tau : |\mu_\tau - \mu_\tau^*(P)| > \epsilon\} \mid y_1, y_2, \ldots, y_n) \to 0 \text{ as } n \to \infty \qquad (2.8)$$

for any $\epsilon > 0$ and any distribution P. We show that, under mild conditions on the prior $\pi(\mu_\tau)$, the generalized posterior (2.6) is consistent in the above sense.

Theorem 1. *Let P be the true distribution of the loss population and $\mu_\tau^* = \mu_\tau^*(P)$ denote the corresponding true τ-expectile, with $\tau \in (0,1)$ fixed. If the prior $\pi(\mu_\tau)$ is continuous and bounded away from zero on any neighborhood of μ_τ^*, then the generalized posterior (2.6) is consistent.*

Proof. For any $\epsilon > 0$, consider an ϵ neighborhood around μ_τ^*, $A_\epsilon = \{\mu_\tau : |\mu_\tau - \mu_\tau^*| \leq \epsilon\}$ and split the sample space of $\mathbf{y} = (y_1, y_2, \ldots, y_n)$ into two disjoint regions:

$\mathcal{Y}_n = \{(y_1,\ldots,y_n) : |\hat{\mu}_\tau - \mu_\tau^*| \leq \epsilon/2\}$ and $\mathcal{Y}_n^c = \{(y_1,\ldots,y_n) : |\hat{\mu}_\tau - \mu_\tau^*| > \epsilon/2\}$,

where $\hat{\mu}_\tau$ denotes the unique minimizer of the empirical discrepancy function $D_n(\mu_\tau)$ in (2.5), which is known to be a consistent estimator for μ_τ, see Newey and Powell (1987) in detail. Obviously, we have

$$\int_{A_\epsilon^c} \pi(\mu_\tau \mid \mathbf{y})\, d\mu_\tau = I_{\mathcal{Y}_n}(\mathbf{y}) \cdot \int_{A_\epsilon^c} \pi(\mu_\tau \mid \mathbf{y})\, d\mu_\tau + I_{\mathcal{Y}_n^c}(\mathbf{y}) \cdot \int_{A_\epsilon^c} \pi(\mu_\tau \mid \mathbf{y})\, d\mu_\tau,$$

where $I_B(x)$ denotes an indicator function on the set B. Consistency of $\hat{\mu}_\tau$ implies that the second term vanishes as n goes to infinity and thus we only need to analyze the posterior probability assuming the observed data sets $\mathbf{y}$ reside in $\mathcal{Y}_n$. Rewrite this posterior probability as a ratio N_n/I_n, where

$$N_n = \int_{A_\epsilon^c} \exp\{-\omega n(D_n(\mu_\tau) - D_n(\mu_\tau^*))\}\pi(\mu_\tau)\, d\mu_\tau$$

and

$$I_n = \int_{-\infty}^{\infty} \exp\{-\omega n(D_n(\mu_\tau) - D_n(\mu_\tau^*))\}\pi(\mu_\tau)\, d\mu_\tau.$$

(a) We will show that for any $a > 0$,

$$I_n \geq \exp\{-wan\} \text{ as } n \to \infty. \tag{2.9}$$

For each $\mu_\tau \in A_\epsilon$, by the law of large numbers, it follows

$$D_n(\mu_\tau) - D_n(\mu_\tau^*) \to D(\mu_\tau) - D(\mu_\tau^*) \text{ as } n \to \infty. \tag{2.10}$$

Also the continuity of $D(\mu_\tau)$ guarantees that

$$|D(\mu_\tau) - D(\mu_\tau^*)| < \epsilon/2 \tag{2.11}$$

for any $\mu_\tau \in A_\delta = \{\mu_\tau : |\mu_\tau - \mu_\tau^*| \le \delta\} \subset A_\epsilon$ when δ is small enough. Combining (2.10) and (2.11), it follows

$$D_n(\mu_\tau) - D_n(\mu_\tau^*) \;>\; -\epsilon/4 \tag{2.12}$$

when n is large enough. Therefore, using Fatou's lemma,

$$\begin{aligned}
\liminf_{n\to\infty} \exp\{n\omega\epsilon\} I_n &\ge \liminf_{n\to\infty} \int_{A_\delta} \exp\{n\omega(\epsilon - D_n(\mu_\tau) + D_n(\mu_\tau^*)))\}\pi(\mu_\tau)\, d\mu_\tau \\
&> K \cdot \int_{A_\delta} \pi(\mu_\tau)\, d\mu_\tau
\end{aligned}$$

because $\exp\{n\omega(\epsilon - D_n(\mu_\tau) + D_n(\mu_\tau^*)))\} \to \infty$ as $n \to \infty$. Therefore, (2.9) holds for any $a > 0$ for large n since ϵ is arbitrary.

(b) Next we show that

$$N_n \;=\; \int_{A_\epsilon^c} \exp\{-\omega n(D_n(\mu_\tau) - D_n(\mu_\tau^*))\}\pi(\mu_\tau)\, d\mu_\tau \le \exp\{-t\omega n\} \tag{2.13}$$

for some $t > 0$.

In the following we assume $D_n(\mu_\tau^* - \epsilon) \le D_n(\mu_\tau)$ and the case when $D_n(\mu_\tau^* - \epsilon) > D_n(\mu_\tau)$ can be considered similarly.

For any $\mu_\tau \in A_\epsilon^c$, it follows

$$\begin{aligned}
D_n(\mu_\tau) - D_n(\mu_\tau^*) &= D_n(\mu_\tau) - D_n(\mu_\tau^* - \epsilon) + D_n(\mu_\tau^* - \epsilon) - D_n(\mu_\tau^*) \\
&\ge D_n(\mu_\tau^* - \epsilon) - D_n(\mu_\tau^*)
\end{aligned}$$

because of the fact that $\hat{\mu}_\tau \notin A_\epsilon^c$ and the convexity of the function $D_n(\cdot)$. Therefore,

$$N_n = \int_{A_\epsilon^c} \exp\{-\omega n(D_n(\mu_\tau) - D_n(\mu_\tau^*))\}\pi(\mu_\tau)\, d\mu_\tau \le \exp\{-\omega n(D_n(\mu_\tau^* - \epsilon) - D_n(\mu_\tau^*))\}.$$

Note that as $n \to \infty$, $D_n(\mu_\tau^* - \epsilon) - D_n(\mu_\tau^*) \to D(\mu_\tau^* - \epsilon) - D(\mu_\tau^*)$ by the law of large numbers and hence, (2.13) holds true due to the continuity of $D(\cdot)$ and the uniqueness of μ_τ^*.

Combining (2.9) and (2.13), we have

$$Pr(\mu_\tau \in A_\epsilon^c \mid y_1, y_2, \ldots, y_n) \;\le\; \frac{N_n}{I_n} \;\le\; \exp\{-\omega(t/2 - a)n\}.$$

Since this holds for any $a > 0$, so taking $a < t/2$ shows that the left-hand side above must converge to 0 as $n \to \infty$, which proves the consistency claim. □

$\square$

2.4 Numerical Analysis

2.4.1 Simulation Study

We present a small simulation study to evaluate the finite sample performance for our proposed approach. Assume that the population distribution of losses, denoted by $Pa(\alpha, \beta)$, is the Pareto distribution with the shape parameter $\alpha > 0$ and the scale parameter $\beta > 0$ whose probability density function is given by

$$f(x; \alpha, \beta) = \frac{\alpha\beta^{\alpha}}{(x+\beta)^{\alpha+1}}, \;\; x > 0.$$

It is well-known that its mean is $\frac{\beta}{\alpha-1}$ if $\alpha > 1$ and its variance is $\beta^2 \frac{\alpha}{(\alpha-1)^2(\alpha-2)}$ if $\alpha > 2$. Direct calculation shows that the equation (2.3) becomes

$$\frac{\tau}{1-\tau} = \left(1 + \frac{\mu_\tau}{\beta}\right)^{\alpha-1} \left[(\alpha-1)\frac{\mu_\tau}{\beta}\right]^{\beta} - 1 \tag{2.14}$$

for given $0 < \tau < 1$. The τ-expectile could then be easily solved for any $\alpha > 1$ and β from this equation. For our simulation we set $\alpha = 3$ and $\beta = 100$. Therefore, the true or theoretical 99%-expectile of the population $Pa(3, 100)$ is $\mu^*_{0.99} = 323.37$. We simulate a dateset from $Pa(3, 100)$ with the sample size $n = 300, 500$ and 1000, respectively. For each dataset, the Bayesian estimate (posterior mean) and the maximum likelihood estimate (MLE) of the $\mu_{0.99}$ are calculated. We repeat the simulation 500 times for each sample size. The average of the Bayesian estimates and the average of the MLEs with corresponding standard deviations in the parentheses are presented in the following table.

Estimates	$n = 300$	$n = 500$	$n = 1000$
MLE	320.35(65.24)	324.59(55.60)	324.42(41.24)
Bayesian estimate	324.35 (116.60)	324.38(97.54)	323.93(70.68)

We can see from the table that both MLE and Bayesian estimate are very close to the true expectile for each case even when the sample size is 300; and as expected, the standard deviation of each estimate decreases as the sample size increases. Note that the standard deviation of the MLE is always less than the corresponding one for Bayesian estimate. This has no surprise because the

MLE was obtained when assuming the population is Pareto distributed while Bayesian estimate was calculated without assuming anything on the population distribution.

2.4.2 Real Data Analysis

We analyze one commercial loss dataset in Asia-Pacific during the period 2000-2013 for illustration. This data set, collected by the Insurance Risk and Finance Research Centre (www.IRFRC.com), contains 526 large commercial risk losses caused by man-made risks, such as fire, explosion etc, which is also available in R package *CASdatasets* (see http://cas.uqam.ca/). The losses range from EUR 140K to 366 million, see the following frequency table for the losses (in millions) in different sub-intervals. For further details, refer to the technical report by Benedetti and Milidonis (2015). Our purpose here is to produce a generalized posterior distribution for the expectile of this unknown loss distribution when $\tau = 0.99$.

losses	(0.14, 1.00]	(1.00, 5.00]	(5.00, 10.00]	(10.00, 50.00]	(50, 100]	(100, 366]
counts	201	100	73	109	25	18

The empirical estimate $\hat{\mu}_{0.99}$ of $\mu_{0.99}$ is 150.35, which can be obtained through minimizing the empirical discrepancy function (2.3.1). We use a $Gamma(a, b)$ prior for the expectile, such that its prior mean is around $\hat{\mu}_{0.99}$ and its standard deviation is $c\hat{\mu}_{0.99}$, where $c \in (0, 1]$. This can be achieved by taking the gamma shape and scale parameters as $a = c^2$ and $b = c^2\hat{\mu}_{0.99}$, respectively. Here we take $c = 2^{-1/2}$, so that $a = 2$ and $b = \hat{\mu}_{0.99}/2 = 75.17$. This prior is only a suggestion and our numerical result shows that the analysis is not sensitive to the choice of the prior. Using the approach of the unit information loss, we obtained the estimated tuning parameter $\hat{\omega} = 0.003$. With these settings, the generalized posterior (2.6) of the expectile $\mu_{0.99}$ is depicted in Figure 1.

The generalized posterior distribution (2.6) contains all the current information about the expectile $\mu_{0.99}$ of the unknown loss distribution. For many practical purposes, we could summarize the posterior inference as desired. For example, the posterior mean and median are both around 150.23. In addition to point summaries, an interval summary can also be easily obtained. A 95% credible interval for the expectile $\mu_{0.99}$ is $(142.65, 157.81)$, which reports the summary of the expectile $\mu_{0.99}$ with some uncertainty.

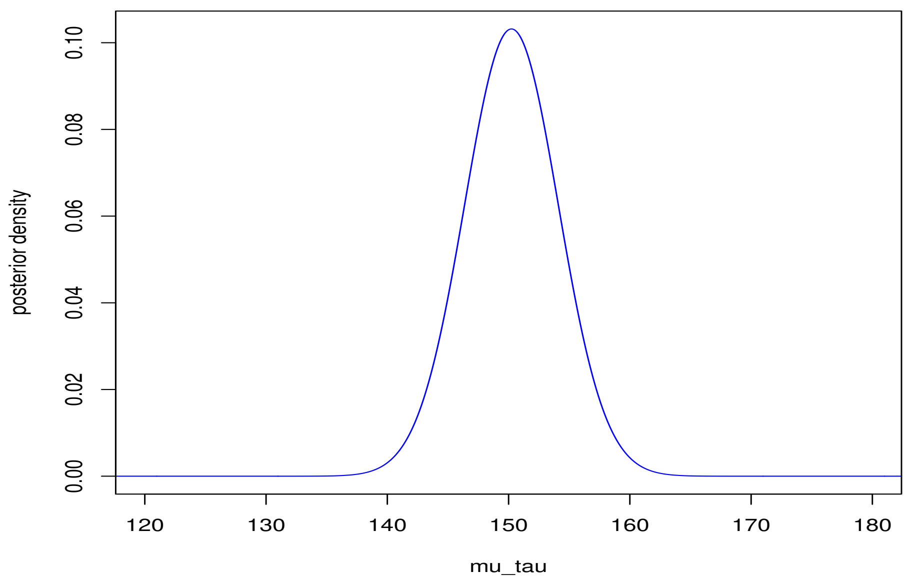

Figure 2.1: The generalized posterior density for the expectile $\mu_{0.99}$

2.5 Concluding Remarks

Expectiles generalize the expectation just as quantiles generalize the median. Quantiles and expectiles both measure the tail behavior of a distribution function although they are in distinct ways. The expectile is the only financial risk measure satisfying coherence and elicitability and is becoming an increasingly popular tool in risk management and capital allocation for financial institutions. Therefore, estimation of the expectile is an important task in financial research. Motivated by Bissiri et al. (2016) and Syring et al. (2019), we have developed a framework for making inference on the population expectile through the generalized Bayesian approach. Our methodology does not need to specify the parametric distribution for the loss population and thus avoids the potential model misspecification issue. It is hoped that the method presented in this article could be extended to other complicated models such as the varying-coefficient model (Xie et al. (2014)) or the conditional autoregressive model (Taylor (2008)) in the future.

References

Antweiler, W., & Frank, M. Z. (2004). Is all that talk just noise? The information content of internet stock message boards. The Journal of finance, 59(3), 1259-1294.

Artzner, P., Delbaen, F., Eber, J.-M. & Heath, D. (1999). Coherent measures of risk. Mathematical Finance 9, 203–228.

Baik, B., & Nam, H. J. (2009). The effect of regulation fair disclosure on conference calls: The case of earnings surprises. Asia-Pacific Journal of Financial Studies, 38(6), 801-829.

Bartov, E., Givoly, D., & Hayn, C. (2002). The rewards to meeting or beating earnings expectations. Journal of accounting and economics, 33(2), 173-204.

Bellini, F. & Bignozzi, V. (2015). On elicitable risk measures. Quantitative Finance 15, 725–733.

Bellini, F. & Di Bernardino, E. (2017). Risk management with expectiles. The European Journal of Finance 23, 487–506.

Benedetti, D., B. E. & Milidonis, A. (2015). Large commercial risks (LCR) in insurance: Focus on asia-pacific. Tech. Rep. 1, Insurance Risk and Finance Research Centre.

Bissiri, P. G., Holmes, C. C. & Walker, S. G. (2016). A general framework for updating belief distributions. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 78, 1103–1130.

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.

Brown, L. D., & Caylor, M. L. (2005). A temporal analysis of quarterly earnings thresholds: Propensities and valuation consequences. The Accounting Review, 80(2), 423-440.

Brunswick. (2014). The investment community's use of digital and social media. Retrieved from http://www.brunswickgroup.com/media/266919/Infographic-Digital-InvestorSurvey.pdf

Chemmanur, T. J., & Tian, X. (2014). Communicating private information to the equity market before a dividend cut: An empirical analysis. Journal of Financial and Quantitative Analysis, 49(5-6), 1167-1199.

Chen, H., De, P., Hu, Y. J., & Hwang, B. H. (2014). Wisdom of crowds: The value of stock opinions transmitted through social media. The Review of Financial Studies, 27(5), 1367-1403.

Chen, J. M. (2018). On exactitude in financial regulation: Value-at-risk, expected shortfall, and expectiles. Risks 6, 1–28.

Daouia, A., Girard, S. & Stupfler, G. (2018). Estimation of tail risk based on extreme expectiles. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 80, 263–292.

Gao, Xing, The Information Content of CEOs' Personal Social Media: Evidence from Stock Returns and Earnings Surprise (May 1, 2019). Working Paper.

He, J., Liu, T., Netter, J., & Shu, T. (2020). Expectation management in mergers and acquisitions. Management Science, 66(3), 1205-1226.

Holzmann, H. & Klar, B. (2016). Expectile asymptotics. Electronic Journal of Statistics 10, 2355–2371.

Hong, L. & Martin, R. (2020). Model misspecification, bayesian versus credibility estimation, and gibbs posteriors. Scandinavian Actuarial Journal 2020, 634–649.

Jacquier, E. & Polson, N. (2011). Bayesian methods in finance. The Oxford Handbook of Bayesian Econometrics 1, 439–512.

Keating, C. & Shadwick, W. F. (2002). A universal performance measure. Journal of Performance Measurement 6, 59–84.

Kratschmer, V. & Zahle, H. (2017). Statistical inference for expectile-based risk measures. Scandinavian Journal of Statistics 44, 425–454.

Kross, W. J., Ro, B. T., & Suk, I. (2011). Consistency in meeting or beating earnings expectations and management earnings forecasts. Journal of Accounting and Economics, 51(1-2), 37-57.

Martin, R. (2014). Expectiles behave as expected. Risk 2014, 79–83.

Matsumoto, D. A. (2002). Management's incentives to avoid negative earnings surprises. The accounting review, 77(3), 483-514.

Miller, G. S., & Skinner, D. J. (2015). The evolving disclosure landscape: How changes in technology, the media, and capital markets are affecting disclosure. Journal of Accounting Research, 53(2), 221-239.

Newey, W. K. & Powell, J. L. (1987). Asymmetric least squares estimation and testing. Econometrica: Journal of the Econometric Society 55, 819–847.

Philipps, C. (2021). Interpreting expectiles. Available at SSRN 3881402 .

Rachev, S. T., Hsu, J. S., Bagasheva, B. S. & Fabozzi, F. J. (2008). Bayesian Methods in Finance. John Wiley & Sons.

Schmeiser, H., Siegel, C. & Wagner, J. (2012). The risk of model misspecification and its impact on solvency measurement in the insurance sector. The Journal of Risk Finance 13, 285–308.

Skinner, D. J., & Sloan, R. G. (2002). Earnings surprises, growth expectations, and stock returns or don't let an earnings torpedo sink your portfolio. Review of accounting studies, 7(2), 289-312.

Syring, N., Hong, L. & Martin, R. (2019). Gibbs posterior inference on value-at-risk. Scan dinavian Actuarial Journal 7, 548–557.

Taylor, J. W. (2008). Estimating value at risk and expected shortfall using expectiles. Journal of Financial Econometrics 6, 231–252.

Tetlock, P. C. (2011). All the news that's fit to reprint: Do investors react to stale information?. The Review of Financial Studies, 24(5), 1481-1512.

Uppal, R. & Wang, T. (2003). Model misspecification and underdiversification. The Journal of Finance 58, 2465–2486.

Xie, S., Zhou, Y. & Wan, A. T. (2014). A varying-coefficient expectile model for estimating value at risk. Journal of Business & Economic Statistics 32, 576–592.

Yamai, Y., Yoshiba, T. et al. (2002). On the validity of value-at-risk: comparative analyses with expected shortfall. Monetary and Economic Studies 20, 57–85.

Ziegel, J. F. (2016). Coherence and elicitability. Mathematical Finance 26, 901–918.