// NO.25-040 | 07/2025

DISCUSSION PAPER

// JOHANNES J. GAUL AND PASCAL SCHRADER

# Beyond Content: Investors' Chatter, Interaction and Earnings Announcement Returns

Leibniz Association

ZEW

SFB/Transregio 266
ACCOUNTING FOR TRANSPARENCY

WORKING PAPER SERIES

No. 207 | July 2025

Gaul, Johannes J. | Schrader, Pascal

# Beyond content
# Investors' chatter, interaction and earnings announcement returns

TRR 266 Accounting for Transparency

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation):
Collaborative Research Center (SFB/TRR) – Project-ID 403041268 – TRR 266 Accounting for Transparency

www.accounting-for-transparency.de

# Beyond Content

## Investors' Chatter, Interaction and Earnings Announcement Returns*

Johannes J. Gaul†, Pascal Schrader‡

July 30, 2025

**Abstract**

We study the relationship between investors' social media activity and earnings announcement returns. To distinguish between information contained in peer-to-peer interaction and user-posted content, we analyze conversation networks on Reddit using centrality metrics from network science and classify user sentiment with large language models. We show that pre-announcement sentiment is positively associated with short-term cumulative abnormal returns only if it does not spark pre-announcement controversy. If pre-announcement controversy arises, we document a negative association. Our findings present a more nuanced view on the wisdom of crowds hypothesis, highlighting that peer-to-peer interaction on social media exhibits a pattern of normalization, and thus contains informational value beyond content.

**Keywords:** Information Processing, Reddit Wallstreet Bets, Wisdom of Crowds, Conversation Networks, Large Language Models, Eigenvector Centrality, High-frequency Data

**JEL classification:** G12, G14

*We are grateful for valuable comments and suggestions from Eli Bartov, Holger Daske and participants at the Accounting Research Camp on Transparency in Corporations and Markets (ARCTIC) 2025, the Humboldt University Berlin Brown Bag Seminar, and the European Accounting Association (EAA) Annual Conference 2025. We acknowledge financial support from the German Research Foundation (DFG) for the project "Accounting for Transparency" (Grant No. SFB/TRR 403041268). Johannes Gaul and Pascal Schrader also thank the Graduate School of Economics and Social Sciences (GESS) at the University of Mannheim for financial support. All remaining errors are our own. The authors declare no conflicts of interest.

†ZEW, University of Mannheim, NeSt (johannes.gaul@zew.de)

‡University of Mannheim (pascal.schrader@uni-mannheim.de)

# 1 Introduction

"Post-truth" refers to a phenomenon in which individuals are guided more by emotions and personal beliefs than by facts (McIntyre, 2018). Recognizing its growing relevance, the Oxford English Dictionary named it the "Word of the Year" in 2016. This designation signals a broader perception that is particularly prevalent in, although not exclusive to, information exchanged on social media. There, false information, a "meme culture", and insulting language find fertile ground (Jimenez-Duran, 2023). In recent years, accounting and finance research investigates the role of information exchanged on social media and recognizes that these platforms are important channels for investors to discuss market activity (e.g., Campbell et al., 2023; Drake et al., 2017; Lerman, 2020). The general finding is that sentiment expressed there about particular securities correlates with their subsequent market returns (Bartov et al., 2018; Berkovitch et al., 2023).

We question how these findings resonate with the "post-truth" phenomenon and aim to shed light on the associated mechanisms. We capture the "post-truth"-culture by measuring peer-to-peer social media interaction and as a complement to previous approaches that consider sentiment *only* when assessing associations between social media activity and subsequent capital market returns. Prior research suggests that individuals are more likely to engage with content that challenges their pre-existing beliefs (Mochon and Schwartz, 2024; Taber and Lodge, 2006) and may even derive satisfaction from confronting sources that provoke anger (Lerner and Tiedens, 2006). As a result, users tend to express disagreement more frequently than agreement. Elevated levels of engagement–such as replies–may thus indicate controversy rather than consensus. If there is a normalizing function the previously documented positive association between expressed sentiment and capital market returns, e.g., Bartov et al. (2018), may even be attenuated.

Given that sentiment in social media discussions is *on average* positively associated with subsequent market performance of respective securities, to what degree is this average observed sentiment the product of a normalization within a conversation network? To address this question, we make use of the fact that users' (investors') engage or do not engage in conversations on social media selectively. Their action is subject to trading-off the benefits of participation and its opportunity costs.[1] Hence, observing participation of a given investor in a social media conversation is informative, as it reveals her preference.

[1]Benefits encompass social, psychological, and monetary incentives. Social incentives include satisfaction derived from group affiliation. Psychological incentives involve the gratification gained from exercising emotional expression, such as venting. Monetary incentives include financial motives, such as advocating for positions in which users have a vested interest, aiming to influence others to purchase and thereby raise the asset price. Opportunity costs are equivalent to the forgone utility from pursuing activities outside the social media domain. Thus, users engage in conversations selectively, responding primarily when their utility from participation outweighs associated costs.

We operationalize by investigating the relationship between pre-earnings announcement controversy and earnings announcement returns. First, we conduct a detailed analysis of the timing of discussions related to earnings announcements, clearly differentiating between pre-announcement discussions and post-announcement reactions by leveraging high-frequency data. Second, we investigate the relationship between investor sentiment and capital market returns. Finally, we distinguish between the content of individual posts and the dynamics of peer-to-peer interaction in the pre-announcement period. To do so, we decompose the user sentiment by considering peer-to-peer interaction and informational content of posts. We measure controversy in the peer-to-peer context by user centrality in the conversation networks and informational content of posts through large language models (LLM).

Our sample consists out of comments gathered from the "Daily Discussion Threat" on the Reddit subreddit "Wallstreetbets" (WSB) in a rolling, daily fashion. We focus on WSB for several reasons. First, with around 13 million users at the time of data collection, it is one of the largest investor forums. Second, at the time access via the the application programming interface (API) provided precise timestamps for each post, facilitating accurate measurement. Third, the comment structure of WSB allows for constructing user conversation networks. Lastly, the setting enables us to observe which stocks are discussed by investors relative to a baseline, i.e. the discussion of all other stocks aside from the focal stock. Because we analyze daily discussion threads on WSB, rather than platforms like X, we can observe both presence and absence of mentions, providing a more comprehensive measure of investor interest. We *do not* claim that WSB comments move the market, but rather that it offers a broad reflection of investors' sentiment and pre-announcement controversy. Altogether, we collect data from public conversations on the forum for 310 trading days. We did this before the platform restricted access to its APIs in July 2023. This approach results in a broad sample of 3,300,043 comments. Our dataset includes the precise timing (second) at which each comment was submitted. We focus our analysis on S&P 500 listed firms, yielding 150,660 firm-day observations, 2,430 of which are on the day of a firm's earnings announcement. We decompose the conversation network by computing user eigenvector centrality measures and use a fine-tuned bidirectional encoder representation from transformers (BERT) model to assess the sentiment of each comment. Lastly, we classify users as "informed" (or not) with the ChatGPT o1-preview model.

We begin by examining whether investors actively communicate mentioning stock tickers in the pre-announcement period. We find that investors show consistent interest in firms' earnings

announcements, as evidenced by their discussions on WSB. In particular, investor chatter increases during the trading day *prior* to the announcement but remains high until around 20 hours afterwards. This observation contributes to the ongoing debate about whether investors respond to information events. For example, Lerman (2020) suggests that investors react to earnings announcements, whereas Drake et al. (2017) finds that less sophisticated investors exhibit lower interest in corporate events, suggesting that heightened attention to anticipated news may take precedence over post-announcement discourse.

Consistent with prior studies, we find that the well-documented association between overall sentiment and capital market returns holds in the context of our sample, when compared to other platforms, such as X, used in Bartov et al. (2018). On average, a one standard deviation change (on earnings-days StDev of sentiment = 6.76) in sentiment on the day of the earnings announcement is reflected in a positive CAR[1,2] of around 27 basis points, when controlling for the earnings surprise. We interpret this finding as sentiment by investors on social media platforms such as Reddit, foreshadowing sentiment of the overall investor community.

Concerning our main research question, we find that the association between user posts and market outcomes is contingent on the degree of peer-to-peer engagement. The posts' content of informed users has a positive association with cumulative abnormal returns (CAR) if it is uncontroversial, but the relationship turns negative when it triggers controversy. Given that conversations evolve endogenously before the earnings announcement date, we claim that controversy itself (i.e. heightened peer-to-peer interaction) contains informational value beyond content. In other words, the identification of a user's content as controversial emerges from the peers' reaction and debate ex-ante the earnings announcement, thereby providing a normalizing function or *service* to the community as a whole.

We contribute to the ongoing discussion on the relationship between social media and capital markets, (e.g., Bartov et al., 2018; Berkovitch et al., 2023; Jia et al., 2020), by providing associative evidence that the informational value embedded in peer-to-peer interaction extends beyond the textual content of posts. We interpret this finding as an additional dimension of wisdom of crowds, which can be predictive of capital market outcomes. By documenting conversation networks prior to earnings announcements, we gain deeper insight into both the "common ground" and the "controversy" within the investor community. Selective participation in (online) discussions follows a normalizing tendency. Furthermore, sentiment expressed by uninformed users in the ex-ante period correlates positively, although less strongly, with subsequent market outcomes.

Our findings have practical implications. Retail trading has risen steadily in recent years (Eaton et al., 2022). Fintech trends such as neobrokers and fractional-share investing have further lowered entry barriers and expanded market participation. This shift has also changed how retail investors exchange information. Platforms such as SeekingAlpha, Estimize, X, and Reddit are becoming increasingly popular for retail investors. Users discuss and contextualize firms' disclosures, such as earnings announcements. Our findings are thus informative not only to the scientific community but also to retail investors, investment professionals and investment firms that rely on social media platforms to gather information for their forecasting models (Brown, 2024). Against this backdrop, it is unsurprising that many social media outlets have begun to privatize API access in recent years. Our findings underscore the importance of monitoring peer-to-peer interaction, as it changes the relationship between sentiment and capital market outcomes.

The remainder of the paper is organized as follows: Section 2 contains a review of the accounting and finance literature and the development of our hypothesis. In section 3 the data sources and measures are defined. In section 4 we introduce the empirical design. We describe the sample in section 5, present and discuss results in section 6. We conclude in section 7.

## 2 Related Literature and Hypothesis

Wisdom of crowds is often cited by accounting and finance scholars as an explanation for the positive relationship between social media interaction and capital market outcomes. It posits that large groups of people collectively possess more accurate insights and make better decisions than individual experts. This phenomenon is grounded in diversity of opinion, independence of thought, and decentralized knowledge among participants (Surowiecki, 2004). One classic example is the observation that judges' collective estimates exceeds the trial by an average single judge and even more advanced weighting schemes (Clemen and Winkler, 1986). Evidence from the accounting literature has shown that firms rely on wisdom of crowds in their corporate governance by using investors' opinions to uncover additional value-relevant information in mergers and acquisitions (Ang et al., 2021) and that the business outlook of employees is informative in predicting future operating performance (Huang et al., 2020).

We relate our research to three strands of the accounting and finance literature that are connected to wisdom of crowds. The first stream examines whether users of social media platforms pay attention to earnings announcements. Establishing this link is important, as variation in user attention is a necessary condition for any potential association with capital market outcomes. The second

strand investigates how social media discussions relate to capital market reactions. We extend this literature by incorporating peer-to-peer interaction as an additional predictor. The third strand explores the role of controversy in social media discussions surrounding earnings announcements. While prior studies typically capture controversy as the divergence in sentiment across posts, we focus on the conditional effect of sentiment in the presence of pre-announcement controversy. This approach allows us to separately analyze the informational value of sentiment and disagreement, rather than conflating the two.

### 2.1 Investor Attention to Earnings Announcements

The evidence from the first strand of the literature–whether social media users are interested in earnings announcements–is mixed. Drake et al. (2017) suggest that non-professional investors are either less interested in or slower to react to corporate events such as earnings announcements. Similarly, Campbell et al. (2023) find that while most earnings announcements receive little attention, a few gain significant traction and go "viral". In contrast, Lerman (2020) documents that investors on Yahoo Finance and StockTwits do react to earnings announcements. Kimbrough et al. (2024) further show that consumers' perceptions of firms shift in response to earnings news, suggesting that these announcements reach a broad audience. Additionally, Drake et al. (2012) provide evidence that investor attention, as measured by Google search activity, increases before, peaks during, and remains elevated after earnings announcements–consistent with gradual information diffusion.

While prior studies examine investor attention at the day-firm level, we extend this literature by exploiting the precise timestamps of Reddit comments, i.e. high-frequency data, and earnings announcement timing, to analyze attention at an intraday level. This granularity allows us to distinguish between discussions occurring *before* and *after* the earnings announcement, providing a more nuanced view of how investors respond to anticipated information events.

### 2.2 Sentiment and Capital Market Reactions

The second strand of the literature examines if social media discussions are associated with capital market outcomes. While prior studies draw on a variety of platforms, we restrict this section of the literature review to research using X, StockTwits, and Reddit as their user communities and posting dynamics more closely resemble those of Reddit compared to platforms like SeekingAlpha or Estimize. Differences in character limits, user demographics, interaction types (e.g., posts, comments, tagging), recommendation algorithms, and levels of anonymity can all influence how attention and

sentiment spread across these platforms.

Berkovitch et al. (2023) find that firms' tweets during quarterly earnings announcements are associated with stronger stock price and volume reactions, reduce investor uncertainty, and improve the timeliness and efficiency with which stock prices reflect new information, thereby reducing the post-earnings announcement drift. Bartov et al. (2018) find that aggregate opinions from individual users' tweets prior to earnings announcements successfully predict a firm's forthcoming earnings and announcement returns, especially in weaker information environments. Curtis et al. (2014) show that increased investor attention, measured through social media activity leads to more sensitive pricing of earnings surprises and reduces mispricing, evidenced by diminished post-earnings announcement drift. Conversely, Campbell et al. (2023) reveal that when earnings announcements go viral on social media it can harm markets by decreasing liquidity and slowing price formation. Booker et al. (2023) observe that investor disagreement correlates with increased trading volume around earnings announcements, mainly when disclosure processing costs are low. However, Lerman (2020) notes that while investors in online communities pay attention to accounting information during earnings releases, this increased attention does not affect the earnings response coefficient.

In recent years, accounting researchers have also investigated the association between information disseminated on Reddit and capital market outcomes. Most of the findings are centered on the *Gamestop short squeeze* in January 2021. The evidence on the role of retail investors' coordination on WSB is mixed. Allen et al. (2025) document a sudden price increase for Gamestop and several S&P 500 companies due to coordination of retail investors and Betzer and Harries (2022) find that Reddit posts on WSB significantly and positively influenced GameStop's trading activity in the short term. Extending the findings to Gamestop, AMC, Blackberry, and Nokia during the 2020–2021 decentralized short squeeze, Lyócsa et al. (2022) document that increased activity on WSB relative to Google search trends can explain part of the next day's price variation. On the contrary, Long et al. (2023) find a limited role of the sentiment of comments on Gamestop returns in the 1-30 min window, similarly to Chacon et al. (2023) providing insights that WSB was associated with increased trading volume during the GameStop short squeeze, but a trading strategy that follows its buy and sell recommendations is not profitable. Furthermore, the Gamestop short squeeze altered the culture of WSB, leading to decreased information quality concerning due diligence recommendations (Bradley et al., 2021).

Exceptions from the Reddit literature focused on Gamestop are studies by Eaton et al. (2022), documenting that stocks recently discussed on WSB are more likely to be purchased by retail

investors on Robinhood. Similarly, Münster et al. (2024) show that Robinhood investors purchase stocks recommended on Reddit, regardless of the sentiment. Semenova and Winkler (2021) and Hu et al. (2023), investigate the effects of different investor beliefs, positions, and social dynamics on asset prices. Hu et al. (2023) show that opinions from influential Reddit users significantly shape the beliefs of naïve investors and predict future returns. Semenova and Winkler (2021) further demonstrate that highly upvoted posts lead to coordinated buying and abnormal returns.

### 2.3 Investor Disagreement on Social Media

The third strand of the literature explores how investor disagreement is related to the market's ability to incorporate information, particularly in the context of social media. For example, difference-of-opinions models challenge the notion of fully rational investors who converge toward a shared belief upon receiving the same information. Instead, these models suggest that when investors hold diverse priors and face uncertainty about signal precision, public disclosures can lead to divergent posterior beliefs. Armstrong et al. (2024) show that investors, uncertain about signal credibility, tend to overweight signals confirming their priors, exacerbating belief polarization. Similarly, Glaeser and Sunstein (2013) describe "asymmetric Bayesianism", where agents rationally dismiss information inconsistent with their pre-existing views, even when signals are balanced. Banerjee (2011) and Zhou and Chen (2018) further argue that relative overconfidence–investors' tendency to overweight their own signal relative to others–reduces reliance on market prices, diminishing price informativeness. These models predict that in settings of high polarization or disagreement, investors selectively process information, leading to weaker price reactions even when sentiment is strong. In such settings, price changes may reflect the average of conflicting updates, while trading volume captures the underlying dispersion in belief updating. This body of work establishes a theoretical foundation for anticipating lower information aggregation in polarized environments, such as social media platforms where users interact selectively and interpret content through belief-consistent filters. We cannot observe different prior beliefs independently from one another and we do not measure posterior beliefs, which would both be required to test the difference-of-opinions hypothesis. However, we make use of the fact that any pre-announcement controversy can be indicative of normalizing behavior on social media.[2] Recent empirical work documents the role of disagreement on social media platforms, particularly around earnings announcements. Booker et al. (2023) use StockTwits

[2]Consider the case of an investor holding a long position of a security. She may have an incentive to share optimistic sentiment about the upcoming events, such as the earnings announcement, in order to influence others. However, if her view is not widely shared and other investors respond critically, the ensuing disagreement could reveal that her opinion is not reflecting the broader sentiment, thereby weakening the predictive power of her post for future returns and *invoking* the wisdom of crowds phenomenon.

data to show that pre-announcement disagreement and changes in disagreement around earnings events are positively associated with trading volume. Their findings suggest that social media–based disagreement captures investor heterogeneity that is not fully reflected in analyst forecasts and that lower disclosure processing costs amplify the effects of disagreement. Similarly, Hu et al. (2023) show that sentiment dispersion on Reddit correlates with higher volatility and weaker predictive power of sentiment for returns.

These studies emphasize the level or change in disagreement but assume a persistent relationship between sentiment and market reaction. Our study differs in three key respects. First, we focus on the conditional effect of sentiment under pre-announcement controversy. Second, we measure controversy not through textual dispersion alone, but through peer-to-peer interaction structures (e.g., eigenvector centrality) that capture interaction dynamics. Third, we specifically analyze how controversy interacts with social media sentiment and moderates the reflection of earnings announcements in prices. This emphasis on controversy as a force that alters the effectiveness of sentiment signals, rather than one that simply coexists with them, adds a new dimension to the social media literature in capital markets.

### 2.4 Hypothesis

To examine how controversy plays a role in social media discussions about earnings announcements–and thereby to investigate whether peer-to-peer interaction in social media networks offers insights beyond the informational content of individual posts, it is necessary to differentiate between the content and the peer-to-peer interaction dimension. Prior studies have analyzed the impact of informed users on capital markets. For instance, Wang et al. (2015) demonstrate that although the general content on platforms like SeekingAlpha and StockTwits has minimal correlation with stock performance, contributions from a subset of high-quality users possess predictive value and can outperform broader market indicators. Moreover, exploiting the variation in user types across platforms, Drake et al. (2017) find that only posts from professional or semi-professional information intermediaries (e.g., SeekingAlpha contributors) significantly predict capital markets, whereas contributions from nonprofessional users tend to be noisy and less informative. Based on the result that the content of posts is associated with capital market outcomes, we distinguish posts disseminating information from uninformative posts.

We operationalize by decomposing the pre-announcement sentiment into the content (informed vs. uninformed) and peer-to-peer interaction (controversial vs. uncontroversial) component. Fol-

lowing the arguments in the literature outlined above, we expect the sentiment of informed and uncontroversial posts to be positively associated with CAR. If there is a normalizing tendency in social interaction, we anticipate a negative relationship between informed but controversial posts. Accordingly, we formulate the following hypothesis:

***H:*** The association between the sentiment of informed Reddit users and short-term CAR is positive when the sentiment is uncontroversial and negative when the sentiment is controversial.

Using peer-to-peer interaction as a measure for pre-announcement controversy is an answer to Bianchi et al. (2023), who emphasize that social network analysis is instrumental in understanding how individual interactions can explain critical outcomes like firm performance and investor behavior, otherwise omitted in analyses that work with individual characteristics only. By applying network analysis to social media interaction, we can analyze patterns that provide insights into how information and sentiment about a firm's upcoming earnings announcement is reflected in the investor community.

# 3 Measures

## 3.1 Attention

We measure attention by counting the number of times a firm's ticker is mentioned in WSB Daily Discussion Thread comments. To identify ticker mentions, we use Python's "reticker" package, which leverages pattern recognition techniques (e.g., regular expressions) to extract stock tickers from text. The package scans input text for excerpts commonly associated with ticker symbols, offering an efficient and precise solution for this task.

One advantage of using "reticker" over dictionary-based approaches is its built-in capability to filter out common false positives, such as acronyms or words that resemble stock symbols (e.g., "a" vs. the ticker symbol "A" for "Agilent Technologies, Inc."). Compared to machine learning-based methods, "reticker" is less prone to biases that arise from selectively identifying tickers which are frequently mentioned in a training dataset, thus minimizing confirmation bias. For example, AAPL (Apple, Inc.) would be identified more precisely than other stocks by machine learning methods, as it occurs more frequently in training data–we aim to avoid this systematic error. Therefore using "reticker" ensures that the extracted ticker mentions provide a good measurement of attention.

### 3.2 Sentiment

One challenge underlying our research design is the estimation of sentiment. In contrast to prior research (e.g., Bartov et al., 2018; Hu et al., 2023), we use a LLM to quantify the opinion expressed in individual WSB comments. Traditional word classification schemes, such as dictionaries and naïve Bayes algorithms, are outdated for sentiment analysis by the more timely approaches (see, e.g., King et al., 2017) as originally they often fail to capture context, sarcasm, and nuanced sentiment. Therefore, we employ two advanced language models, Twitter-roBERTa-base and Finbert, which offer enhanced capabilities in understanding context and sentiment in text.

The first approach, Twitter-roBERTa-base, is based on the Roberta (robustly optimized BERT approach) architecture. This transformer-based model has been pre-trained on a corpus of approximately 58 million tweets, enabling it to "understand" complex linguistic patterns and contextual relationships specific to social media text. When applied to Reddit posts, Twitter-roBERTa-base considers the entire content of each post, allowing it to classify the sentiment as negative, positive, or neutral. This approach helps to capture the general tone and sentiment nuances that dictionary-based methods might miss, making it suited for the informal language often found on Reddit.

The second approach employs the Finbert model, specifically fine-tuned for financial sentiment analysis. Finbert is based on the BERT architecture but is trained on financial texts, making it adept at "understanding" the context and sentiment specific to financial language. This model analyzes the words within each tweet, detecting specific negative or positive terms relevant to financial contexts. Finbert's specialized training allows it to identify sentiment with a focus on financial terminology and context, which is particularly beneficial when analyzing market-related tweets.

We manually assess the performance of both models by labeling the sentiment of 1,000 comments by hand. The roBERTa model achieves a considerably higher subjective accuracy than the Finbert model. This is likely because the language used on Reddit closely resembles the language found on platforms like X (formerly Twitter) rather than the more formal and technical language seen in financial texts such as annual reports. Therefore, we utilize the sentiment classification results from the roBERTa model for our analyses. After classifying each post as positive, neutral, or negative, we assign a score of 1, 0, and -1, respectively.

### 3.3 Informedness

We evaluate different measures to proxy for informed users. Even though we end up using a LLM for our preferred specification, we start by classifying informed comments using the Loughran and McDonald (2011) dictionary. The dictionary contains the source of the words. Most are from the "12dicts" word lists, a standard collection of typical English words. This word list is extended by terms occurring in 10-Ks but not in the original dictionary. We assume that informed users draw upon a specific vocabulary and, therefore, classify comments containing one or more 10-K terms as "informed".

We utilize the uncased DistilBERT model for the second measure, which was fine-tuned on the Multi-Genre Natural Language Inference dataset.[3] Fine-tuning equips the model with a broad "understanding" of language nuances across various genres, enabling it to perform zero-shot classification tasks. This means the model can categorize text data into predefined classes without having been explicitly trained on those specific categories. In our context, this allows for the automated classification of comments into relevant categories such as "finance" and "accounting".

Lastly, we use the OpenAI API to classify Reddit comments. By implementing a batch processing method, we input comments into the ChatGPT o1-preview model using a tailored prompt that directs the model to identify comments discussing accounting and finance topics, micro and macroeconomic developments, or SEC regulations. To enhance the model's accuracy, we provide five explicit examples of both comment types in the prompt. The examples includes informed comments discussing company earnings, economic indicators, regulatory development and uninformed comments characterized by speculative hype as well as lack of substantive content. The model provides classifications in a structured JSON format, indicating whether each comment is deemed informed or uninformed. The prompt is provided in appendix A.

We manually evaluate the quality of all three measures based on a random sample and ultimately select the classification provided by ChatGPT's o1-preview model, as it consistently delivers the highest-quality results. Figure 1 provides an illustrative example of informed and uninformed conversation interaction.

### 3.4 Controversy

To measure controversy, we rely on the eigenvector centrality. This metric captures the centrality of a user within the network of Reddit interactions on a given day. For this purpose we construct

[3] See also: `https://huggingface.co/typeform/distilbert-base-uncased-mnli`.

a directed network that has the peers as "nodes" and their answers upon each other as edges. Representing the resulting conversation network as an adjacency matrix $A$, the eigenvector centrality is calculated recursively by taking into account the centrality of a node's predecessor nodes. The centrality for node $i$ is then the i-th element of a left eigenvector associated with the eigenvalue $\lambda$ of a positive maximum modulus. The eigenvector ($\vec{x}$) is defined as: $\lambda \vec{x}^T = \vec{x}^T A$, as described in Berman and Plemmons (1994). For the calculation we use the Python library "NetworkX".

The eigenvector centrality is conceptually similar to Google's PageRank algorithm. It considers not only how often a user is engaged with but also the centrality of those engaging with them. For example, being frequently replied to by highly central users increases a user's eigenvector centrality more significantly than receiving numerous replies from less impactful users. This characteristic is particularly valuable in social media contexts, where the structure of the community and the reputation of interacting users shape the diffusion of ideas. Another advantage of eigenvector centrality in a directed network is that it cannot be inflated artificially through a user's own activity. A user cannot achieve high centrality merely by commenting frequently; instead, they need to attract engagement from substantially engaged others. We provide an example for the calculation on a given day in figure 2. Figure 2a shows the entire conversation network, figure 2b picks up the specific example given in section 3.3 and figure 1.

We interpret high centrality as a proxy for controversy, motivated by prior research showing that users are more likely to engage with content that challenges their beliefs or provokes emotional reactions (Lerner and Tiedens, 2006; Mochon and Schwartz, 2024; Taber and Lodge, 2006). In social media settings, disagreement tends to prompt more interaction than consensus, meaning that elevated engagement is more likely to reflect dissent than endorsement. Thus, high eigenvector centrality indicates that a user sparks reactions. In this way, the structure of the network allows us to detect instances where peer-to-peer interaction exerts a normalizing effect. We define a user as controversial on a given day if their eigenvector centrality falls within the top decile of the day's eigenvector centrality distribution. We classify all comments by that user on a given day as controversial.

# 4 Empirical Design

## 4.1 Are investors attentive to earnings announcements?

It is ex-ante unclear if investors react to earnings announcements and, if so, how attention to and sentiment about a specific stock varies relative to non-earnings announcement dates. On the one

hand, earnings announcements represent a shock to the information environment, and earnings news could be a focal point for discussion among peers. On the other hand, Reddit users are uninformed at large and might be late or less interested in covering these events but may favor the discussion of concurrent topics idiosyncratic to individual stocks instead (e.g., investors might be more interested in Elon Musk taking over Twitter, Inc. than earnings announcements).

To test whether investors on WSB pay attention to earnings announcements, we test if the discussion of a specific stock is centered around the earnings announcement date. We estimate the following model using OLS:

$$Attention_{i,d} = \alpha_0 + \alpha_1 EAD_{i,d} + \alpha_2 Abs(Return_{i,d}) + \alpha_3 (EAD_{i,d} \times Abs(Return_{i,d})) + \phi_i + \lambda_d + \epsilon_{i,d}, \quad (1)$$

where $Attention_{i,d}$, is attention to a given stock, measured as the number of posts containing the ticker of stock $i$ on the day $d$. The test variable $EAD_{i,d}$ is a binary indicator, taking the value of one if firm $i$ announced its earnings on day $d$. $Abs(Return_{i,d})$ is the absolute daily return of the stock. $\phi_i$ and $\lambda_d$ are ticker and day fixed effects, respectively. Controlling for both ticker and day fixed effects is important, as attention on a given stock may be systematically higher (e.g., "meme" stocks, Fisch (2022)) or lower due to conflicting events that co-occur on the same day (e.g., Russia invading Ukraine). We cluster standard errors on the ticker level. In specification 1, the hypothesis that the earnings announcements increase investors' attention to the respective firm implies $\alpha_1 > 0$. As attention is measured as count data, we complement this analysis with a poisson pseudo-maximum likelihood specification taking in the fixed effects as specified.

Additionally, we test whether the sentiment on the earnings announcement day differs from the sentiment on regular dates, even though we have no ex-ante expectation of the direction of the relationship between the earnings announcement date and average sentiment. To test the relationship, we estimate the following regression:

$$Sentiment_{i,d} = \alpha_0 + \alpha_1 EAD_{i,d} + \alpha_2 Return_{i,d} + \alpha_3 (EAD_{i,d} \times Return_{i,d}) + \phi_i + \lambda_d + \epsilon_{i,d}, \quad (2)$$

where $Sentiment_{i,d}$ measures the users' expression. We classify each comment as "positive", "negative," or "neutral" based on the Twitter-roBERTa-base LLM and calculate the sum on the firm-day level. In contrast to specification 1, we refrain from taking the absolute value of returns or the surprise but measure them, including negative values. The reason is that sentiment is not restricted to a minimum measurement of zero compared to our measure for attention. In specification

2, the hypothesis that sentiment differs on earnings announcement dates relative to other dates implies $\alpha_1 \neq 0$. Again, we cluster standard errors on the level of the individual ticker.

### 4.2 Does investor chatter on the earnings announcement day predict CAR?

Before testing our main hypothesis concerning decomposition of sentiment into controversial and uncontroversial, we establish the baseline relationship between sentiment and CAR to align with prior studies. We limit our sample to earnings announcement dates and estimate the following regression:

$$\begin{aligned} CAR[1;N]_{i,d} = \alpha_0 + \alpha_1 Sentiment_{i,d} + \alpha_2 Surprise_{i,d} \\ + \alpha_3 (Sentiment_{i,d} \times Surprise_{i,d}) + \lambda_d + \epsilon_{i,d}, \end{aligned} \tag{3}$$

where the dependent variable $CAR[1;N]_{i,d}$ is the CAR in a window of one to four days after the earnings announcement, calculated using the well-established event study methodology in MacKinlay (1997). We include day fixed effects, because at the time of measurement the relative returns between tickers and relative sentiments of the respective day are known. We provide additional results with the inclusion of other fixed effects structures.

### 4.3 Is peer-to-peer interaction associated with earnings announcement returns?

To answer our research question, we decompose ($Sentiment_{i,d}$) into different categories and attribute it to user characteristics. Figure 1 illustrates this using an exemplary conversation thread. In particular, we investigate whether the sentiment of controversial comments (outlined in red) differs from that of uncontroversial comments. To reflect the heterogeneous nature of user-generated content, we also distinguish between comments that convey substantive information ("informed") and those that do not ("uninformed"), indicated by the dashed blue boxes. Based on Pedersen (2022), we group all users at the daily level into four distinct categories based on two criteria: whether they are controversial or not (determined by their relative ranking of eigenvector centrality) and whether they are informed or not (determined by the content of their comments). After this classification, we assign the previously calculated comment $Sentiment_{i,d}$ to each user group and aggregate it on the ticker-day level. In other words, for each ticker, we determine the aggregate sentiment exposure to each of the four user groups from the two-by-two matrix of informed and/or controversial users comments. Our analysis thus facilitates that users–ex ante to the earnings announcement–react to each other. We estimate the following regression, where $Surprise_{i,d}$ is measured as

previously defined. The main independent variables are the sentiment measures. Specifically, $SentInfoContro_{i,d}$ reflects the aggregate sentiment expressed by informed and controversial users, while $Sent\overline{Info}Contro_{i,d}$ captures sentiment from users who are controversial but not informed. The variable $SentInfo\overline{Contro}_{i,d}$ summarizes sentiment from informed but non-controversial users, and $Sent\overline{InfoContro}_{i,d}$ captures sentiment from users who are neither informed nor controversial.

$$\begin{aligned} CAR[1;N]_{i,d} = &\alpha_0 + \alpha_1 Surprise_{i,d} + \alpha_2 SentInfoContro_{i,d} + \alpha_3 Sent\overline{Info}Contro_{i,d} \\ &+ \alpha_4 SentInfo\overline{Contro}_{i,d} + \alpha_5 Sent\overline{InfoContro}_{i,d} \\ &+ \alpha_6 SentInfoContro_{i,d} \times Surprise_{i,d} + \alpha_7 Sent\overline{Info}Contro_{i,d} \times Surprise_{i,d} \\ &+ \alpha_8 SentInfo\overline{Contro}_{i,d} \times Surprise_{i,d} + \alpha_9 Sent\overline{InfoContro}_{i,d} \times Surprise_{i,d} \\ &+ \phi_i + \epsilon_{i,d}, \end{aligned} \tag{4}$$

These variables allow us to disentangle the role content and peer dynamics play in capital market responses. If there is normalizing behavior in the social media setting, we would expect a negative coefficient on $\alpha_2$, indicating that sentiment expressed by informed-controversial users is negatively associated with CAR in the short-term window following the earnings announcement.

# 5 Data and Sample Selection

## 5.1 Data

To answer our research questions, we collect data from several sources. First, we use the LSEG Workspace API (formerly Thomson Reuters Refinitiv) for stock market data. LSEG offers firm-level data, including stock prices and market capitalization. Second, earnings announcement dates, analyst forecasts, and earnings information are collected from the Institutional Brokers' Estimate System (I/B/E/S). The database compiles the forecasts of various brokerage firms and analysts. It allows researchers to compare analysts' consensus estimates with actual company performance, which can be used to calculate earnings surprises. Third, we use Reddit to collect data about stock discussions. Using the Reddit APIs "pmaw" and "praw", we gather daily data from 310 trading days, spanning from 06/04/2022 to 02/07/2023.[4] We are limited to this sample period, as Reddit constrained access to its API after that date. Reddit is a social media platform that allows users to interact through discussion and visual content. Posts and comments appear under a given username, usually a pseudonym, so users do not appear with their clear names. Comments are

[4]On the 07.04.2022 there was no submission of the Daily Discussion Thread on the platform.

posted on so-called "Subreddits" specific to a certain topic. For any given Subreddit, the number of subscribers is displayed, and the liveliness of the community is visible by the number of daily posts ("submissions") and comments under these posts.[5]

We turn to the widely known Subreddit WSB to gather information on peers' discussions. This Subreddit has been used by other authors contributing to the finance and accounting literature (e.g., Allen et al., 2025; Bradley et al., 2021; Eaton et al., 2022; Hu et al., 2023; Long et al., 2023; Semenova and Winkler, 2021). Specifically, we extract comments and multidimensional peer reactions from the submission titled "Daily Discussion Thread", which is posted on weekdays on WSB. This thread serves as a central hub for regular discussions on current market topics.

Although WSB also features dedicated submissions during earnings seasons–focused on earnings week discussions only–we intentionally exclude this data from our analysis. Instead, we concentrate on the broader discussions within the Daily Discussion Thread. This approach enables us to capture a mix of discussions, including those directly related to earnings announcements and ticker mentions, as well as general comments. By doing so, we establish a baseline for the level of general discourse occurring at the time. Our methodology distinguishes us from studies that selectively search for ticker mentions using "cashtags"[6] on platforms like X or StockTwits. This broader focus allows us to measure the dynamics of overall peer discussions, rather than limiting our scope to conversations explicitly flagged as related to specific stocks.

The submission "Daily Discussion Thread" is a good sampling ground for analyzing the connection between investors' discussion on earnings announcements and price responsiveness. It fosters interaction among individuals with shared interest in the stock market while avoiding a narrow focus on earnings news, encouraging broader discussion. The thread also accommodates conversations on cross-sectional topics affecting the stock market, such as macroeconomic events. Our sampling period includes pivotal events like Elon Musk's acquisition of Twitter, Inc., sharp interest rate increases, and inflationary pressures linked to Russia's invasion of Ukraine. Reddit's low participation barriers–stemming from its ease of use and anonymous nature–enable a diverse group of participants to engage in discussions. While some comments offer potentially valuable insights for investment decisions (e.g., accurate predictions of interest rates, sales forecasts, or economic developments), others consist of boilerplate content, misleading claims, or derogatory remarks.

Low oversight, high anonymity, and a threaded comment structure contribute to a distinct social media environment on Reddit. These factors partly lead to using offensive language and

[5] An illustrative example of such an interaction is provided in figure 1.

[6] A cashtag is a symbol used in online discussions to indicate a specific stock ticker, functioning similarly to hashtags (#), but preceded by a "$" sign. For example, $AAPL refers to the ticker for Apple, Inc.

promoting a "meme culture". This contrasts the more analytical and data-driven discussions on Seeking Alpha, where investors generally favor in-depth, research-oriented posts. It also differs from X, where investors frequently engage in quick, real-time market commentary, with posts reflecting a broader range of expertise levels.

### 5.2 Sample Selection

We start our analysis with all firms listed in the S&P500 (503). We exclude firms without quarterly earnings announcements (11) and missing earnings information (6). We collect comment data in a rolling daily fashion from Reddit using the available APIs for all 310 trading days. We, thus, base our analysis on 150,660 firm-day observations (486 firms * 310 days). Table 2, panel A details the sample selection on the firm level. We plot the average daily returns of our sample firms compared to all S&P 500 firms in appendix figure A.1 to validate that our sample selection does not introduce systematic biases.

On the comment level, we start with all comments posted in the "Daily Discussion Thread" of Reddit's WSB forum. We are able to gather 3,300,043 comments from 59,933 users in total. 817,966 of these comments by 37,754 unique users mention one or multiple stocks. As our sample is restricted to S&P 500 companies, we filter for comments mentioning firms listed in the index (288,154) and remove in the last step firms with missing financials. We are left with 278,116 comments by 19,533 users. Table 2, panel B details the sample selection on the ticker level, resulting in 486 sample tickers (96.6% coverage of all S&P 500 firms).

### 5.3 Descriptive Statistics

Table 3 presents descriptive statistics for Reddit comments and the key capital market variables structured into three panels based on the level of aggregation.

Panel A of table 3 reports individual comment-level characteristics. The comments have an average length of 12.56 words, and users write, on average, a comment every 5.7 days.

Panel B of table 3 provides day-level characteristics. On each trading day (310 observations), on average, 959 ticker symbols of 82.47 firm's are mentioned in comments. This implies that 17% of firms in our sample are discussed on a given day.

Panel C of table 3 displays ticker-day level characteristics based on 150,660 firm-day observations. On average, a ticker is mentioned 1.43 times per day, with the aggregate sentiment across comments being slightly negative (-0.25). The panel further breaks down Reddit activity

by peer-to-peer comment types. On average, comments that are uninformed-controversial (0.52) and uninformed-uncontroversial comments (0.48) occur more often compared to informed comments. Informed-controversial comments appear less often, averaging 0.33 per ticker-day, while informed-uncontroversial comments are the least frequent, averaging 0.09 per ticker-day. Regarding sentiment, uninformed-controversial comments exhibit the most negative average sentiment (-0.09). Informed-controversial comments have a slightly negative average sentiment (-0.07), whereas informed-uncontroversial comments are close to neutral (-0.02).

Finally, panel C of table 3 also shows the capital market variables. The variables used in our main tests, abnormal return and earnings surprise, display effect sizes similar to those in the related literature (e.g., Bartov et al., 2018; Hirshleifer and Sheng, 2022), but negative. This is plausible as our sampling period contains the aftermath of the COVID-19 pandemic and the Russo-Ukraine war. As our observation period spans 310 days, we document approximately five earnings announcements for each firm. This corresponds to 0.2 % of the 150,660 firm-day observations that capture earnings announcements. Furthermore, we display descriptive statistics of variables not directly used in empirical analyses but for calculating the earnings surprise (stock price) and abnormal returns (market capitalization and S&P 500 market return).

In figures 3a and 3b, we provide anecdotal evidence for the heterogeneity in investor interest in firms by visualizing the number of posts mentioning Tesla, Inc. and Walmart, Inc. Investors active on Reddit almost exclusively discuss Walmart, Inc. surrounding the earnings announcement. The only exception is a spike on 25/07/2022 when Walmart, Inc. issued a profit warning. Tesla, Inc. gained more interest from investors on Reddit, even on average trading days. Nonetheless, we capture a similar trend. The investor chatter spikes around the earnings announcements and other major news events idiosyncratic to the firm. For Tesla, Inc., these news, among others, are the removal from the S&P 500 ESG on 18/05/2022, a lawsuit filed by Twitter, Inc. against Elon Musk (20/09/2022), and the announcement of a production stop in the largest factory in Shanghai (17/12/2022). A similar pattern emerges in figure 4 when plotting the earnings announcements and the sum of all comments discussing firms in our sample. On average, the discussions centering around stock tickers increase when firms announce their earnings.

In table 4 we compare the attention to earnings announcements across industries. The results suggest that firms operating in communication services, consumer discretionary, and information technologies see the largest increase in Reddit comments on the firm's respective earnings announcement day. The "Big Tech" firms Alphabet, Inc. (communication services), amazon.com, Inc.

(consumer discretion), and Apple, Inc. (information technology) are active in these industries and highlight anecdotally the interest in their earnings announcements by investors.

Finally, we leverage the exact time stamps of each Reddit WSB comment under the "Daily Discussion Thread" to develop a better understanding of when users are active on Reddit. Figure 5 displays the total number of Reddit posts per minute. Users seem to be most active in the morning between 9 am and 10 am ET and in the afternoon between 3 pm and 4 pm ET. Even though Reddit is used internationally, most discussions take place during the daytime in the United States and during trading hours of the New York Stock Exchange.

# 6 Results

## 6.1 Are investors attentive to earnings announcements?

Table 5 presents the regression results examining the correlation between earnings announcement days (EAD) and investor attention and sentiment. Stock returns (Return) are included as control variables to account for market performance influences on investor behavior. All estimations contain ticker and day fixed effects. Standard errors are clustered on the ticker level.

The dependent variable in columns 1 to 3 is investor attention. Column 1 displays the results without controlling for returns. The coefficient for EAD is positive and statistically significant at the 1% level, indicating that, on average, there are approximately 5.77 more comments posted concerning the specific ticker on earnings announcement days than on non-announcement days. This suggests heightened investor attention during these events. Column 2 displays the relationship between absolute stock returns and Reddit attention. The positive coefficient suggests that Reddit users are more active when absolute returns increase. In column 3, the absolute value of stock returns (Abs(Return)) is added to EAD to control for the influence of market movements on investor attention. The EAD coefficient decreases but remains positive and significant. The coefficient for Abs(Return) is also positive and significant, suggesting that larger stock price movements are associated with increased investor attention. The results indicate a significant increase in attention on earnings announcement days relative to a, on average, low baseline on non-announcement days, after accounting for ticker and day fixed effects.

Columns 4 to 6 focus on investor sentiment as the dependent variable. Column 4 shows the baseline regression without controlling for stock returns. The EAD coefficient is negative and significant at the 1% level, indicating that, on average, one more negative comment is posted in anticipation of the earnings announcement. This could reflect expected disappointment or

heightened scrutiny during these periods. Results in column 5 suggest, that with increasing returns the sentiment on Reddit is more positive. Column 6 incorporates stock returns (Return) to control for the effect of market performance on sentiment. The EAD coefficient remains negative and significant, while the Return coefficient is positive and significant. As a robustness check, we estimate the relationship between attention and earnings announcements using Poisson pseudo-maximum likelihood (PPML) regression, suitable for count data, as presented in appendix table B.1. While our main results employ OLS estimation, the PPML results yield similar patterns, reinforcing the robustness of our findings.

We extend prior literature by utilizing the precise timestamps from Reddit comments to examine user posting behavior within a 48-hour window surrounding earnings announcements. Figure 6 illustrates that investor discussions on Reddit are particularly concentrated in the five hours immediately preceding earnings announcements. This observation is related to prior literature investigating investor interest in earnings announcements. While some studies (e.g., Drake et al. (2012); Lerman (2020)) suggest that investors exhibit interest around earnings announcements, others (e.g., Drake et al. (2017); Gomez et al. (2024)) argue that social media interactions do not increase during these periods. Our descriptive evidence adds to this by documenting that investors actively engage on Reddit during the earnings announcement day. Their discussions reflect both, expectations of and reactions to earnings announcements.

To answer the question whether attention and sentiment capture information beyond what is contained in the earnings surprise, we restrict our sample to the earnings announcement days only. Table 6 presents the regression results analyzing the relationship between investor attention, sentiment, and earnings surprises. The dependent variable in columns 1 to 3 is investor attention. All models incorporate ticker and day fixed effects, with standard errors clustered at the ticker level. Column 1 assesses the relationship between the absolute value of stock returns (Abs(Return)) and attention. The coefficient for Abs(Return) is positive and statistically significant at the 5% level, suggesting that larger stock price movements, irrespective of direction, are associated with increased investor attention. Column 2 examines the impact of the absolute value of the earnings surprise (Abs(Surprise)) on attention. The coefficient is positive but not statistically significant, implying that the earnings surprise alone does not predict investor attention when controlling for ticker and day fixed effects. In column 3, both Abs(Return) and Abs(Surprise) are included as explanatory variables. The coefficient for Abs(Return) remains positive and significant with 65.84, while the coefficient for Abs(Surprise) decreases and remains insignificant.

The dependent variable in columns 4 to 6 is investor sentiment. Column 4 explores the relationship between stock returns (Return) and sentiment. The coefficient for Return is 12.22 and significant at the 1% level, indicating that positive stock returns are associated with more favorable investor sentiment. This aligns with a positive correlation between investors sentiment and market movement, on average. Column 5 investigates the earnings surprise (Surprise) association with sentiment. The coefficient of 11.93 is positive but not statistically significant, suggesting that the earnings surprise does not significantly impact investor sentiment. In column 6, both Return and Surprise are included in the model. The coefficient for Return is 12.15 and remains positive and significant, while the coefficient for Surprise diminishes considerably and remains insignificant.

The lack of significant coefficients for Abs(Surprise) and Surprise when controlling for Returns implies that they capture similar channels within day, whereas attention and sentiment capture separate concepts that are not (multi-)collinear to earnings surprises. This distinction is important, as it establishes that attention and sentiment are not merely proxies for earnings surprises, but can provide additional insights into investor behavior.

### 6.2 Does investor chatter on the earnings announcement day predict CAR?

Table 7 presents the regression results examining whether investor attention and sentiment on Reddit during earnings announcement days is associated with subsequent CAR. In this section, we present the regression results examining the relationship between Reddit user sentiment and CAR following earnings announcements. We consider two day [1,2], three day [1,3], and four day [1,4] post-earnings announcement windows.[7] Columns 1 to 3 display results without ticker fixed effects, while columns 4 to 6 include day and ticker fixed effects. In columns 1 and 4, where the dependent variable is CAR[1,2], the coefficient for Sentiment is positive and statistically significant at the 5% or 1% level. One additional positive comment is, on average, associated with a 4 (12) basis point increase in the CAR two-days after the earnings announcement without (with) ticker fixed effects. This suggests that investor sentiment captured from social media discussions is associated with subsequent short-term market returns, consistent with findings of prior literature.

The coefficient for Surprise is also positive and significant at the 1% level in both columns. This highlights that positive earnings surprises contribute substantially to higher CAR, reinforcing the importance of fundamental financial information in driving stock returns. The interaction term between Sentiment and Surprise is negative, but not significant. This suggests that the positive effect of stock price reaction to more extreme earnings surprises is not changed by social media

[7]We exclude the earnings announcement day [0,1], we also try longer periods, but the results remain unchanged.

chatter.

In columns 2 and 5, extending the CAR window to three days, the coefficient for Sentiment decreases in magnitude and loses statistical significance in column 2, but regains significance at the 5% level in column 5 when ticker fixed effects are included. This indicates that the predictive power of sentiment on CAR weakens over a more extended post-announcement period without fixed effects but remains significant when controlling for day and ticker-specific factors (therefore using information not available during the time of the earnings announcement for estimation). The coefficients for Surprise remain positive and significant at the 1% level, confirming the persistent impact of earnings surprises on stock returns over multiple days. The interaction terms in these columns are negative, but not statistically significant. Similarly, in columns 3 and 6, where the CAR window is extended to four days, the coefficient for Sentiment only becomes significant when including ticker fixed effects.

Overall, the results support the hypothesis that Reddit user sentiment is positively associated with CAR following earnings announcements, particularly in the immediate short term. The significant positive association in the two-day window suggests that sentiment derived from Reddit contains information anticipative of short-term market reactions. However, the relationship weakens over longer horizons. Notably, the results remain significant when including both day and ticker fixed effects. Yet the interpretation changes: The ticker fixed effects are computed using the full return series, which includes future market reactions. As a result, the demeaning process may inadvertently incorporate information not available at the time, potentially inflating the observed association between sentiment and returns.

The generally insignificant interaction effects between sentiment and earnings surprises suggest that the predictive power of social media sentiment does not significantly vary with the magnitude of the earnings surprise. Overall, the results indicate that market participants can benefit from monitoring Reddit discussions to gain insights into short-term stock performance.

### 6.3 Is peer-to-peer interaction associated with earnings announcement returns?

Table 8 reports the regression results for our main research question, analyzing how the relationship between sentiment and CAR changes depending on the level of controversy. CAR is measured over event windows ranging from two to four days post-announcement. Earnings surprise, is consistently positive and statistically significant at the 1% level across all models, with coefficients ranging from 78.18 to 111.96. This confirms the role of fundamental earnings information in explaining market

reactions, as positive earnings surprises are associated with greater abnormal returns.

The results suggest that distinguishing the content (informed vs. uninformed) from the peer-to-peer interaction (controversial vs. uncontroversial) offers additional insights into the relationship between social media sentiment and capital market reactions. While an increase of one additional positive comment by uninformed-uncontroversial users is associated with an increase of 0.19 percentage points in CAR[1,2] after the earnings announcement, this relationship flips if comments spark controversy. One additional positive informed-controversial comment is associated, on average, with a negative CAR[1,2] of -0.1 percentage points. In terms of economic significance, a one standard deviation increase (StDev sentiment uninformed-uncontroversial on earnings day: 1.66) in sentiment of uninformed-uncontroversial users is associated with a 0.32 percentage point increase in CAR[1,2] after the earnings announcement. In contrast, a one standard deviation increase (StDev sentiment informed-controversial on earnings day: 2.94) in sentiment of informed-controversial users is associated with a 0.29 percentage point decrease in CAR[1,2] after the earnings announcement. In the long run, we do not document a statistically significant relationship.

When incorporating ticker fixed effects into the regression models, as presented in columns 4-6 in table 8, we observe some differences in the relationship between user sentiment and CARs. With a coefficient of 0.32, the number of uninformed-uncontroversial comments now exhibit a significant robust positive association at the 1% level with CAR[1,2] and remains significant in the three-day window. In contrast, the number of informed-controversial comments no longer shows a significant relationship with CAR (although still negative), differing from the earlier results without ticker fixed effects where a significant negative association was observed. A noteworthy change is seen in the interaction terms. The interaction between earnings surprise and the number of informed, controversial comments becomes negative and significant at the 1% level in the two-day window and remains significant at the 5% and 10% levels in subsequent windows up to five days. This suggests that while earnings surprises generally are associated with higher CARs, this relationship is attenuated when the sentiment of comments that are informed-uncontroversial is high.

Overall, the results emphasize that decomposing the content and peer-to-peer interaction is useful to learn more about subsequent association with market returns. The association of sentiment is largely transient, with measurable differences diminishing beyond short-term windows. These findings contribute to understanding how differences in opinions, measured by peer-to-peer interaction on social media, complement the relationship between investor sentiment and capital market responses to earnings information.

# 7 Conclusion

This study advances the understanding of the interplay between investor communication on social media and financial markets, particularly in the context of earnings announcements. By leveraging granular data from Reddit's WSB forum and employing LLMs and network analysis, we highlight the dynamics of retail investor chatter and price responsiveness to earnings information.

First, we establish that investor attention in our sample is focused on earnings announcements. We provide descriptive evidence on the intraday level of investor chatter around the earnings announcement time and document heightened attention. Sentiment is more negative on earnings announcement days, although this could be a reflection of the sample period being a time of economic and social distress. Generally, sentiment is strongly correlated with the market return on a given day.

Second, our findings highlight the ability to use social media sentiment to predict earnings announcement returns. We find a positive and significant association between pre-announcement sentiment and CARs. This result contributes to the growing body of literature on earnings announcement trading signals (Lyle and Yohn, 2024) and the effect of social media interaction on price formation. We extend prior findings to the uniquely unfiltered and communal nature of WSB discussions and propose a novel methodology relying on conversation decomposition.

Finally, we examine the normalizing role of controversy in social media interactions. Recognizing that in conversations, users react endogenously to posts, we advance the measurement of sentiment by decomposing it into controversial and uncontroversial components prior to the earnings announcement. This approach not only refines the predictive power of aggregate sentiment–the "wisdom of crowds"–for market reactions, but also allows us to identify controversial posts based on users' revealed preferences to engage. We find that sentiment expressed by controversial users is negatively associated with subsequent market outcomes, suggesting a normalization mechanism: When positive sentiment sparks disagreement, the corrective reactions of other users reduce its influence. Overall, our findings show that peer-to-peer interaction on Reddit is anticipative of CAR and helps uncover the dynamics of collective information processing in retail investor networks.

Taken together, this study provides evidence of the predictive value of investor chatter and peer-to-peer interaction on market outcomes around earnings announcements. Interaction rather than sentiment-only is reflective of a "post-truth" online culture. If content itself is becoming more and more meaningless while opportunity costs to participating in (online) discussions prevail, participation and interaction foreshadow market outcomes.

## References


Allen, F., Haas, M., Nowak, E., Pirovano, M., and Tengulov, A. (2025). Squeezing shorts through social media platforms. *Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3823151*. Cited on pages 6 and 16.

Ang, J. S., Hsu, C., Tang, D., and Wu, C. (2021). The role of social media in corporate governance. *The Accounting Review*, 96(2):1–32. Cited on page 4.

Armstrong, C. S., Heinle, M. S., and Luneva, I. (2024). Financial information and diverging beliefs. *Review of Accounting Studies*, 29(3):2082–2124. Cited on page 7.

Banerjee, S. (2011). Learning from prices and the dispersion in beliefs. *The Review of Financial Studies*, 24(9):3025–3068. Cited on page 7.

Bartov, E., Faurel, L., and Mohanram, P. S. (2018). Can Twitter help predict firm-level earnings and stock returns? *The Accounting Review*, 93(3):25–57. Cited on pages 1, 3, 6, 10, and 18.

Berkovitch, J., Israeli, D., and Subramanian, V. (2023). Firms' Tweets and stock price discovery. *Available at SSRN: https://ssrn.com/abstract=4587965*. Cited on pages 1, 3, and 6.

Berman, A. and Plemmons, R. J. (1994). *Nonnegative Matrices in the Mathematical Sciences*. Classics in Applied Mathematics. SIAM. Cited on page 12.

Betzer, A. and Harries, J. P. (2022). How online discussion board activity affects stock trading: The case of GameStop. *Financial Markets and Portfolio Management*, 36(4):443–472. Cited on page 6.

Bianchi, P. A., Causholli, M., Minutti-Meza, M., and Sulcaj, V. (2023). Social networks analysis in accounting and finance. *Contemporary Accounting Research*, 40(1):577–623. Cited on page 9.

Booker, A., Curtis, A., and Richardson, V. J. (2023). Investor disagreement, disclosure processing costs, and trading volume evidence from social media. *The Accounting Review*, 98(1):109–137. Cited on pages 6 and 7.

Bradley, D., Hanousek Jr, J., Jame, R., and Xiao, Z. (2021). Place your bets? The market consequences of investment research on Reddit's Wallstreetbets. *Available at SSRN: https://ssrn.com/abstract=3806065*. Cited on pages 6 and 16.

Brown, M. (2024). What AI sees in the market (that you might not) – ten ways investors are, or should be, using large language models. *Chicago Booth Review*, Fall 2024:28–39. Cited on page 4.

Campbell, B., Drake, M., Thornock, J., and Twedt, B. (2023). Earnings virality. *Journal of Accounting and Economics*, 75(1):101517. Cited on pages 1, 5, and 6.

Chacon, R. G., Morillon, T. G., and Wang, R. (2023). Will the reddit rebellion take you to the

moon? evidence from wallstreetbets. *Financial Markets and Portfolio Management*, 37(1):1–25. Cited on page 6.

Clemen, R. T. and Winkler, R. L. (1986). Combining economic forecasts. *Journal of Business & Economic Statistics*, 4(1):39–46. Cited on page 4.

Correia, S., Guimaraes, P., and Zylkin, T. (2020). Fast poisson estimation with high-dimensional fixed effects. *Stata Journal*, 20(1):95–115. Cited on page 45.

Curtis, A., Richardson, V. J., and Schmardebeck, R. (2014). Investor attention and the pricing of earnings news. *Available at SSRN: https://ssrn.com/abstract=2467243*. Cited on page 6.

Drake, M. S., Roulstone, D. T., and Thornock, J. R. (2012). Investor information demand: Evidence from Google searches around earnings announcements. *Journal of Accounting Research*, 50(4):1001–1040. Cited on pages 5 and 20.

Drake, M. S., Thornock, J. R., and Twedt, B. J. (2017). The internet as an information intermediary. *Review of Accounting Studies*, 22(2):543–576. Cited on pages 1, 3, 5, 8, and 20.

Eaton, G. W., Green, T. C., Roseman, B. S., and Wu, Y. (2022). Retail trader sophistication and stock market quality: Evidence from brokerage outages. *Journal of Financial Economics*, 146(2):502–528. Cited on pages 4, 6, and 16.

Fisch, J. E. (2022). Gamestop and the reemergence of the retail investor. *Boston University Law Review*, 102:1799. Cited on page 13.

Glaeser, E. L. and Sunstein, C. R. (2013). Why does balanced news produce unbalanced views? *Available at SSRN: https://ssrn.com/abstract_id=2230996*. Cited on page 7.

Gomez, E. A., Heflin, F., Moon, J. R., and Warren, J. D. (2024). Financial analysis on social media and disclosure processing costs: Evidence from Seeking Alpha. *The Accounting Review*, 99(5):223–246. Cited on page 20.

Hirshleifer, D. and Sheng, J. (2022). Macro news and micro news: complements or substitutes? *Journal of Financial Economics*, 145(3):1006–1024. Cited on page 18.

Hu, D., Jones, C. M., Zhang, V., and Zhang, X. (2023). The rise of reddit: How social media affects belief formation and price discovery. *Available at SSRN: ssrn.com/abstract=3807655*. Cited on pages 7, 8, 10, and 16.

Huang, K., Li, M., and Markov, S. (2020). What do employees know? Evidence from a social media platform. *The Accounting Review*, 95(2):199–226. Cited on page 4.

Jia, W., Redigolo, G., Shu, S., and Zhao, J. (2020). Can social media distort price discovery?

evidence from merger rumors. *Journal of Accounting and Economics*, 70(1):101334. Cited on page 3.

Jimenez-Duran, R. (2023). The economics of content moderation: Theory and experimental evidence from hate speech on twitter. *Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4590147*. Cited on page 1.

Kimbrough, M. D., Paharia, N., Wang, X., and Wei, S. (2024). The brand value of earnings: An event study of consumer responses to earnings announcements. *The Accounting Review*, 99(3):259–285. Cited on page 5.

King, G., Lam, P., and Roberts, M. E. (2017). Computer-assisted keyword and document set discovery from unstructured text. *American Journal of Political Science*, 61(4):971–988. Cited on page 10.

Lerman, A. (2020). Individual investors' attention to accounting information: Evidence from online financial communities. *Contemporary Accounting Research*, 37(4):2020–2057. Cited on pages 1, 3, 5, 6, and 20.

Lerner, J. S. and Tiedens, L. Z. (2006). Portrait of the angry decision maker: How appraisal tendencies shape anger's influence on cognition. *Journal of Behavioral Decision Making*, 19(2):115–137. Cited on pages 1 and 12.

Long, S., Lucey, B., Xie, Y., and Yarovaya, L. (2023). "I just like the stock": The role of Reddit sentiment in the GameStop share rally. *Financial Review*, 58(1):19–37. Cited on pages 6 and 16.

Loughran, T. and McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, 66(1):35–65. Cited on page 11.

Lyle, M. R. and Yohn, T. (2024). Measuring portfolio gains: The case of earnings announcement trading signals. *The Accounting Review*, 99(4):315–338. Cited on page 24.

Lyócsa, Š., Baumöhl, E., and Vỳrost, T. (2022). Yolo trading: Riding with the herd during the gamestop episode. *Finance Research Letters*, 46:102359. Cited on page 6.

MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of Economic Literature*, 35(1):13–39. Cited on page 14.

McIntyre, L. C. (2018). *Post-Truth*. MIT Press, Cambridge, Massachusetts. Cited on page 1.

Mochon, D. and Schwartz, J. (2024). The confrontation effect: When users engage more with ideology-inconsistent content online. *Organizational Behavior and Human Decision Processes*, 185:104366. Cited on pages 1 and 12.

Münster, M., Reichenbach, F., and Walther, M. (2024). Robinhood, reddit, and the news: The

impact of traditional and social media on retail investor trading. *Journal of Financial Markets*, 71:100929. Cited on page 7.

Pedersen, L. H. (2022). Game on: Social networks and markets. *Journal of Financial Economics*, 146(3):1097–1119. Cited on page 14.

Semenova, V. and Winkler, J. (2021). Social contagion and asset prices: Reddit's self-organised bull runs. *Available at: https://arxiv.org/abs/2104.01847*. Cited on pages 7 and 16.

Surowiecki, J. (2004). *The Wisdom of Crowds: Why the Many Are Smarter Than the Few, and How Collective Wisdom Shapes Business, Economies, Societies, and Nations*. Random House Large Print, New York, 1st large print ed. edition. Cited on page 4.

Taber, C. S. and Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, 50(3):755–769. Cited on pages 1 and 12.

Wang, G., Wang, T., Wang, B., Sambasivan, D., Zhang, Z., Zheng, H., and Zhao, B. Y. (2015). Crowds on wall street: Extracting value from social investing platforms. *CSCW '15: Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 17–30. Cited on page 8.

Zhou, Y. and Chen, R. (2018). Rational expectations, difference of opinions and asset pricing. *Applied Economics*, 50(31):3331–3337. Cited on page 7.

# 8 Figures and Tables

## 8.1 Figures

Figure 1: Example for a conversation on Reddit's Daily Discussion Thread

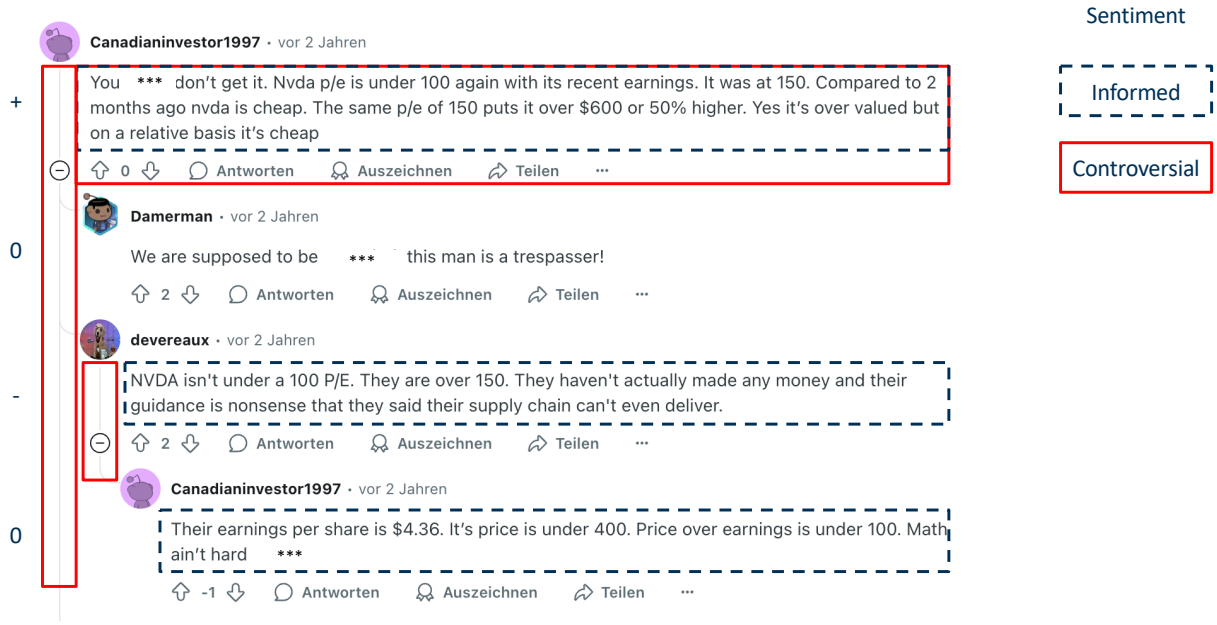


*Note:* Screenshot from the "Daily Discussion Thread for May 25, 2023" on Reddit. Users are arguing about the earnings announcement of NVIDIA. Comments are visually annotated for illustrative purposes: **Blue dashed boxes** indicate *informed* comments, **red boxes** indicate controversial comments based on peer-to-peer interaction. (+), (−), and (0) highlight the sentiment of comments. **Disclaimer:** The language and views expressed in this Reddit conversation are solely those of the individual commenters and do not reflect the opinions or endorsement of the authors. This content is provided for illustrative purposes only.

Figure 2: Conversation network on May 25, 2023

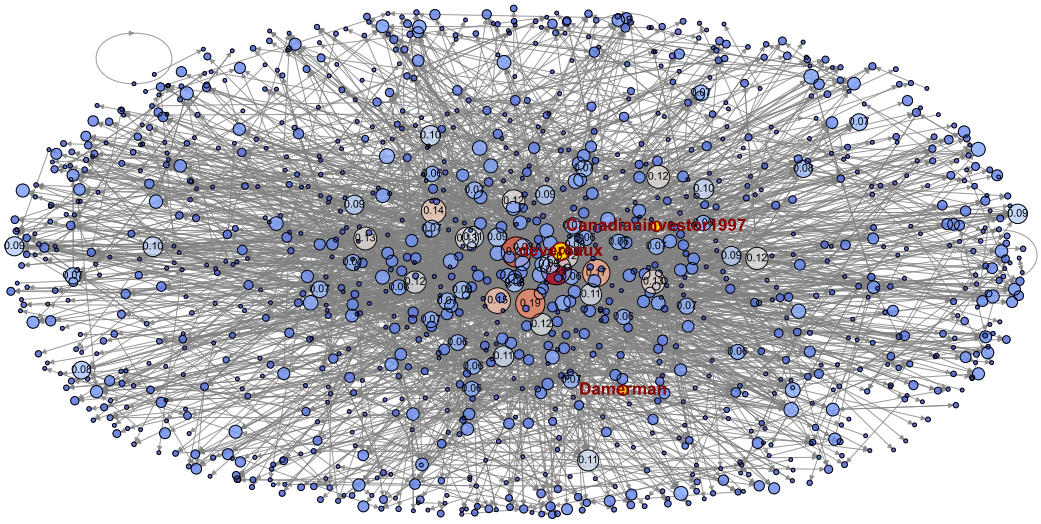


(a) Directed network graph of users engaging in conversation

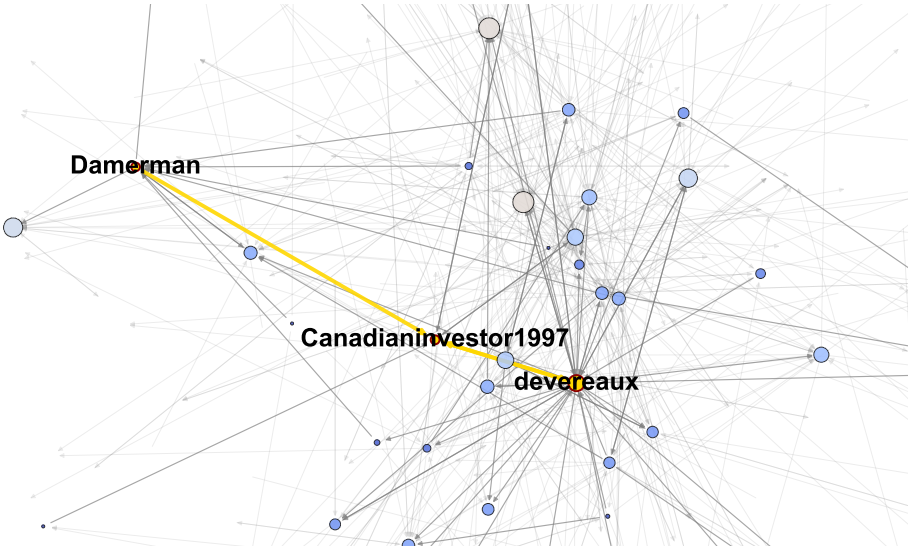


(b) Excerpt of users engaging in conversation

*Notes:* This figure shows the directed conversation network under the Reddit WSB Daily Discussion Thread on May 25, 2023. Figure 2a shows the entire network graph. Users are visualized as nodes with sizes that correspond to their eigenvector centralities. Edges indicate if they answer to each others posts. The arrow is in the direction of the comment, e.g. the arrow points to the node that received the reply. Figure 2b highlights the example given in figure 1 and contains a topologically equivalent excerpt of figure 2a. The conversation excerpt from above is highlighted by the yellow edges. Furthermore included are nodes that have a direct linkage to the three nodes in question. Third degree connections are drawn in a lighter color and third degree nodes are excluded for a better visibility.

Figure 3: Walmart, Inc. and Tesla, Inc. earnings announcements and Reddit discussion

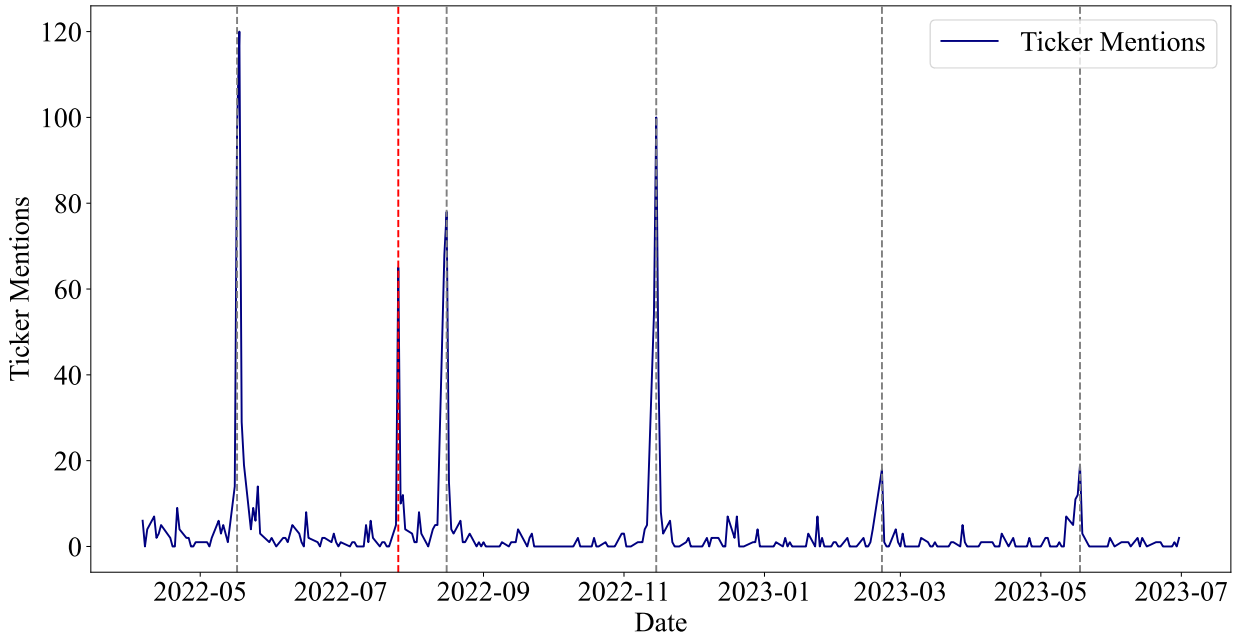


(a) Walmart, Inc. earnings announcements and Reddit discussion

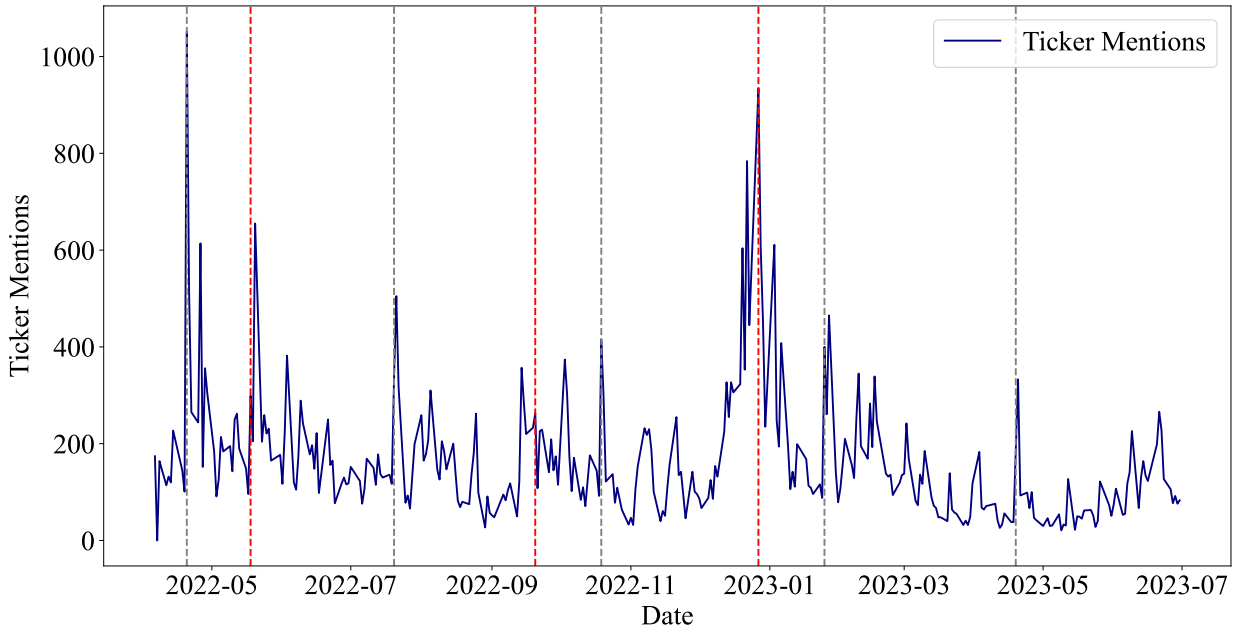


(b) Tesla, Inc. earnings announcements and Reddit discussion

*Notes:* This figures gives a graphical display of the number of Reddit comments about Walmart, Inc. (figure 3a) and Tesla, Inc. (figure 3b). The vertical gray lines indicate the earnings announcement days. The red lines indicate other major corporate events. For Walmart, Inc., on 2022-07-25 a profit warning was disclosed. For Tesla, Inc., on 2022-05-18 it was removed from the S&P 500 ESG, on 2022-09-20 Twitter, Inc. filed a lawsuit against Elon Musk, on 2022-12-27 Tesla, Inc. announces production stop in its larges factory in Shanghai.

Figure 4: Earnings Announcement and Reddit discussion

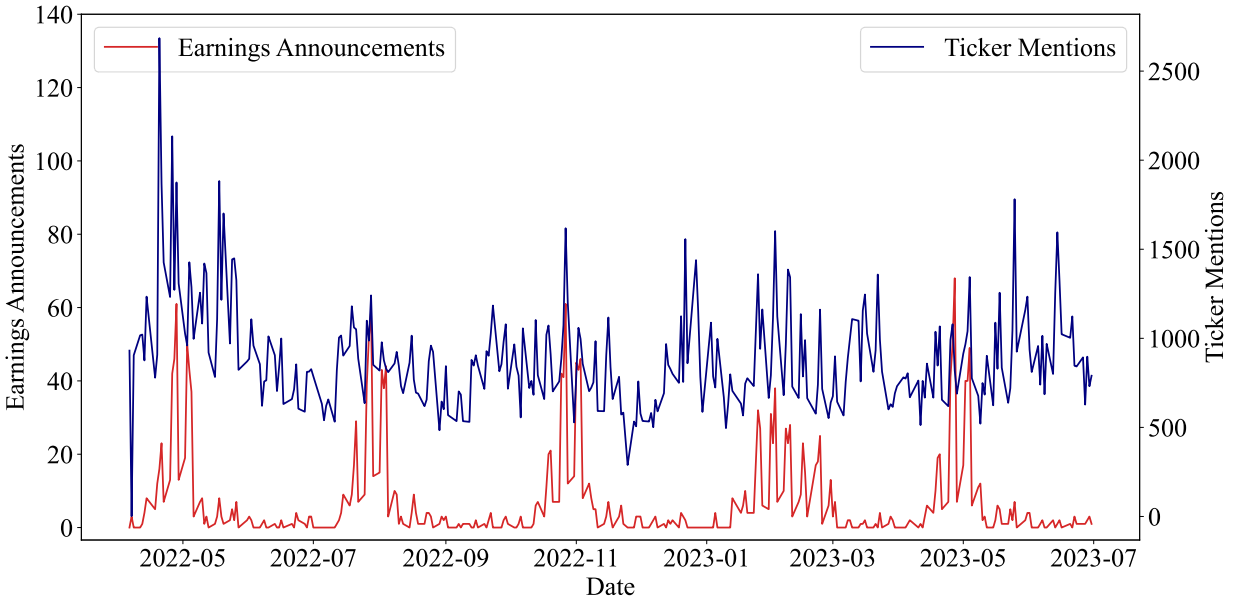


*Notes:* Graphical display of the number of Reddit comments and earnings announcements over the whole sample period of 310 trading days. The red line indicates the number of earnings announcements on a given day (left scale), the blue line the number of comments containing a ticker symbol (right scale).

Figure 5: Traffic on Reddit Wallstreet Bets "Daily Discussion Thread"

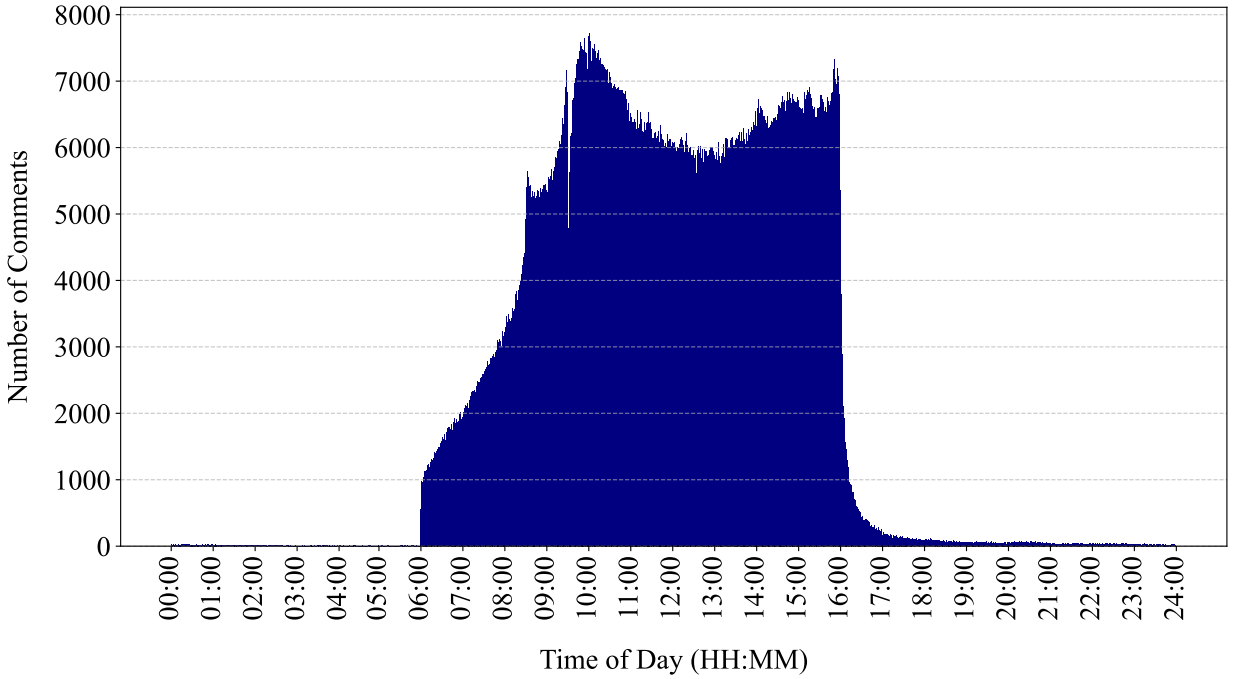


*Notes:* The figure presents a histogram of the number of comments received on a given minute of the day over the whole sample period of 310 trading days. The figure displays New York City local time. The "Daily Discussion Thread" opens each day at 6 am. Previously nobody can comment on the thread, because it does not exist on Reddit. At 6am people start commenting. However, it is clearly visible that the propensity to comment on the "Daily Discussion Thread" decreases drastically with market closing at 4pm Eastern (NY local time), when Reddit posts the daily threat "What are your Moves Tomorrow?".

Figure 6: Total Number of Comments around the Earnings Announcement

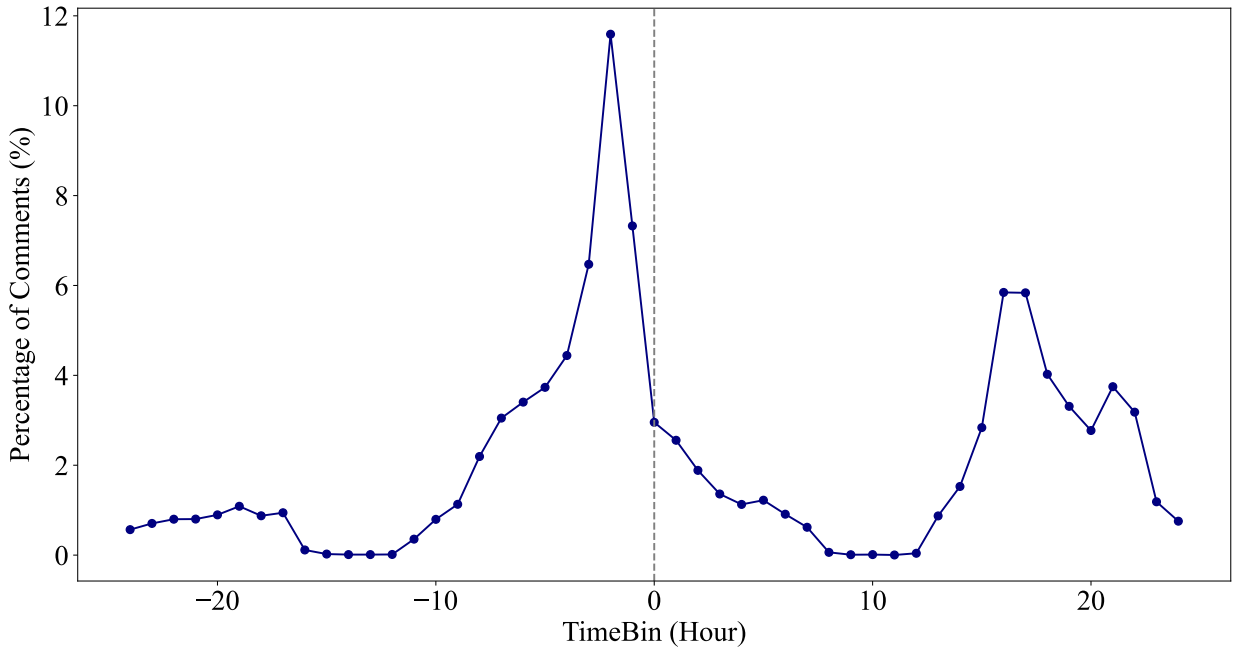


*Notes:* The figure presents the relative number of comments received in the 48-hour window around the earnings announcement. Comments are grouped on the hour level and aggregated for each ticker. Note that for some stocks the earnings announcements take place after market closing. At this point in time the "WSB Daily Discussion Threat" is still online, however rarely used. The second peak, at around 15-18 hours post earnings announcement, shows that attention on the given stocks is higher also after the earnings announcement.

Table 1: Variable Descriptions

| **Variable** | **Definition** | **Source** |
|---|---|---|
| Attention | Attention is the total number of posts in the WSB forum on Reddit mentioning the stock ticker. | Reddit, own calculation. |
| CAR | Abnormal returns are calculated by using event study methodology (MacKinlay, 1997) using the equation $AR_{i,t} = r_{i,t} - (\alpha + \beta \times R_{m,t})$, where $r(i,t)$ is the return on security $i$ in period $t$, and $R(m,t)$ is the market return in period $t$. $\alpha$ and $\beta$ are model parameters, with $\alpha$ assumed to be zero, and $\beta$ calculated by regressing stock returns against market returns. The CAR is then calculated as $CAR_{[t,T]} = \sum_{t=\tau+h}^{\tau+H} AR_{i,t}$, where $T$ is the number of days in the event window. | Earnings announcement dates and stock price data used for the calculation are from I/B/E/S and LSEG respectively. |
| EAD | EAD is an indicator variable equal to 1 if the firm announced earnings on that day, and 0 otherwise. | I/B/E/S |
| Controversial | A user is classified as controversial, if his eigenvector centrality is in the upper decile of the probability density function on a given day. Sentiment of all comments by controversial users are aggregated to the firm-day level. | Reddit, own calculation. |
| Informed | A user is classified as informed if one comment on a given day is classified as informed, according to the Loughran and McDonald (2011) dictionary or ChatGPTs o1-preview model. Sentiment of all comments by informed users are aggregated to the firm-day level. | Reddit, own calculation. |
| Return | Return is calculated as the ratio of the close price of stock $i$ on day $t$ minus the close price on day $t_{-1}$. | LSEG (formerly Refinitiv) |
| Sentiment | Sentiment is the sum of all positive and negative posts on WSB for stock $i$ on day $t$. Positive and negative posts are classified using a RoBERTa model. | Reddit, own calculation. |
| Surprise | Surprise is the deviation between real earnings and the latest consensus analyst forecast divided by the stock price. | I/B/E/S |

*Notes:* This table presents the description of variables used in the empirical analysis and the respective databases.

Table 2: Sample Selection Criteria

| **Panel A: Firm-level** | | |
|---|---|---|
| **Criteria** | **# Firm-Days** | **# Unique Firms** |
| Firms listed in S&P 500 | 156,550 | 503 |
| Firms with quarterly earnings announcements | 152,520 | 492 |
| Firms with daily trading data | 150,660 | 486 |
| **Panel B: Comment Level** | | |
| **Criteria** | **# Comments** | **# Users** |
| All comments in WSB | 3,300,043 | 59,933 |
| Comments mentioning Stock Ticker | 817,966 | 37,754 |
| Comments mentioning S&P 500 Ticker | 288,154 | 19,832 |
| Comments mentioning tickers in sample | 278,116 | 19,533 |

*Notes:* This table presents sample selection criteria. We scraped comments daily from 06/04/2022 to 02/07/2023 (310 trading days). We therefore base our analysis on 150,660 ticker-day observations (486 firms * 310 trading days).

Table 3: Descriptive Statistics

| | Obs. | Mean | St. Dev. | Min | Max |
|---|---|---|---|---|---|
| ***Panel A: Comment-Level Characteristics*** | | | | | |
| Number of words per comment | 3,300,043 | 12.56 | 13.20 | 1 | 1,398 |
| Number of comments per user | 59,933 | 54.82 | 370.07 | 1 | 66,942 |
| ***Panel B: Day-Level Characteristics*** | | | | | |
| Number of tickers / day | 310 | 959.33 | 357.88 | 3 | 2,729 |
| Number of unique tickers / day | 310 | 82.47 | 12.33 | 3 | 122 |
| Market Return / Day [%] | 310 | 0.31 | 1.38 | -4.35 | 5.76 |
| Market Intraday Return / Day [%] | 310 | 0.35 | 1.11 | -3.52 | 4.54 |
| ***Panel C: Ticker-Day Level Characteristics*** | | | | | |
| Number of comments per ticker and day | 150,660 | 1.43 | 12.41 | 0 | 977 |
| Sentiment per ticker and day | 150,660 | -0.25 | 3.08 | -222 | 23 |
| *Peer-to-Peer Comment Types:* | | | | | |
| Informed & Controversial | 150,660 | 0.33 | 3.10 | 0 | 330 |
| Informed & Uncontroversial | 150,660 | 0.09 | 0.84 | 0 | 59 |
| Uninformed & Controversial | 150,660 | 0.52 | 4.88 | 0 | 379 |
| Uninformed & Uncontroversial | 150,660 | 0.48 | 4.24 | 0 | 303 |
| *Sentiment by Comment Type:* | | | | | |
| Informed & Controversial | 150,660 | -0.07 | 0.87 | -91 | 9 |
| Informed & Uncontroversial | 150,660 | -0.02 | 0.31 | -16 | 6 |
| Uninformed & Controversial | 150,660 | -0.09 | 1.08 | -99 | 16 |
| Uninformed & Uncontroversial | 150,660 | -0.06 | 0.83 | -59 | 17 |
| *Stock Characteristics:* | | | | | |
| Market Capitalization [Billion $] | 150,660 | 74.94 | 190.88 | 1.79 | 3,050.90 |
| Price (Open) [$] | 150,660 | 162.58 | 286.11 | 1.77 | 6,344.80 |
| Abnormal Return [%] | 150,660 | -0.03 | 1.38 | -5.74 | 4.34 |
| Abnormal Intraday Return [%] | 150,660 | 0.04 | 1.11 | -3.51 | 4.51 |
| Surprise [%] | 150,660 | -0.08 | 1.46 | -19.07 | 18.33 |
| Earnings Announcement (Binary) | 150,660 | 0.02 | 0.13 | 0 | 1 |

*Notes:* This table presents summary statistics for our main analyses. The sample spans 310 trading days. Variables are grouped by their granularity: individual comment-level metrics, daily market-level indicators (310 observations), and ticker-day specific peer interaction and market data (150,660 observations).

Table 4: Investor Attention and Sentiment by Industry

| | Attention | | | | Sentiment | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) Const. | (2) Abs(Ret.) | (3) EAD | (4) Int. | (5) Const. | (6) Ret. | (7) EAD | (8) Int. |
| Communication Services | -1.47 | 221.19 | 21.20* | -141.39 | -0.41*** | 34.07 | -3.54 | -35.60 |
| N = 5,890 | (2.74) | (152.64) | (11.97) | (138.12) | (0.05) | (24.73) | (2.33) | (43.53) |
| Consumer Discretionary | 1.04 | 135.17 | 15.95* | -129.89 | -0.69*** | 21.94 | -2.05* | -9.11 |
| N = 16,740 | (1.89) | (98.11) | (8.15) | (105.48) | (0.01) | (16.51) | (1.04) | (8.90) |
| Consumer Staples | 0.56*** | 10.47 | -3.75 | 260.83 | -0.12*** | 2.90 | -1.37** | 36.53 |
| N = 11,160 | (0.17) | (15.88) | (5.69) | (197.09) | (0.01) | (2.10) | (0.66) | (32.64) |
| Energy | 0.18 | 11.81* | 2.14* | 19.40 | -0.02*** | 2.48* | 0.15* | 15.58*** |
| N = 7,130 | (0.13) | (5.98) | (1.05) | (24.24) | (0.00) | (1.44) | (0.08) | (4.31) |
| Financials | 0.62*** | 14.84** | 3.46*** | -5.33 | -0.12*** | 2.76 | -0.78*** | -2.07 |
| N = 21,390 | (0.12) | (7.40) | (1.31) | (24.21) | (0.01) | (1.85) | (0.29) | (5.29) |
| Health Care | 0.96*** | 10.20*** | 0.40 | -3.51 | -0.26*** | 1.08** | 0.04 | -0.04 |
| N = 18,600 | (0.05) | (3.54) | (0.24) | (3.14) | (0.00) | (0.45) | (0.11) | (0.82) |
| Industrials | 0.04 | 16.21** | 0.80 | 39.56 | -0.05*** | 2.05*** | -0.65*** | 17.22* |
| N = 22,630 | (0.11) | (6.93) | (1.15) | (44.74) | (0.00) | (0.78) | (0.24) | (10.11) |
| Information Technology | 0.55 | 137.50* | 17.35** | -141.27 | -0.45*** | 4.65 | -1.89** | 6.80 |
| N = 19,840 | (1.40) | (72.75) | (6.98) | (92.91) | (0.01) | (2.99) | (0.77) | (9.85) |
| Materials | 0.48*** | 1.54* | 0.03 | 0.42 | -0.05*** | 0.70 | 0.06 | 1.25 |
| N = 8,680 | (0.02) | (0.90) | (0.13) | (2.64) | (0.00) | (0.53) | (0.04) | (1.42) |
| Real Estate | 0.38*** | 0.42 | -0.07 | 2.24 | -0.11*** | 0.44 | 0.01 | -1.22 |
| N = 9,610 | (0.01) | (0.47) | (0.22) | (5.15) | (0.00) | (0.32) | (0.11) | (1.81) |
| Utilities | 0.60*** | 0.17 | -0.22* | 11.79* | -0.08*** | 0.03 | 0.07 | -7.72* |
| N = 8,990 | (0.02) | (1.44) | (0.11) | (5.92) | (0.00) | (0.43) | (0.06) | (3.84) |

*Notes:* This table presents regression results for specifications 1 and 2 estimated on a sample restricted on the respective industry (rows). Columns 1 to 4 show results with attention as the dependent variable, columns 5 to 8 contain results with respect to sentiment as the dependent variable. The industry classification contains all companies in the S&P500 according to their GICS classification. Each line is a separate regression, one with attention as a dependent variable, one with sentiment as the dependent variable. The regressions are estimated with standard errors clustered on the ticker level, with day and ticker fixed effects. Standard errors are given in parentheses.

Table 5: Earnings Announcements, Investor Attention, and Sentiment

| | Attention | | | Sentiment | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| EAD | 5.77*** | | 5.13*** | -0.93*** | | -0.92*** |
| | (1.16) | | (1.60) | (0.20) | | (0.20) |
| Abs(Return) | | 66.74*** | 61.55*** | | | |
| | | (20.47) | (21.21) | | | |
| EAD×Abs(Return) | | | -6.35 | | | |
| | | | (36.96) | | | |
| Return | | | | | 7.66*** | 7.19** |
| | | | | | (2.86) | (3.00) |
| EAD×Return | | | | | | 4.80 |
| | | | | | | (5.68) |
| Constant | 1.34*** | 0.39 | 0.39 | -0.23*** | -0.25*** | -0.23*** |
| | (0.02) | (0.32) | (0.34) | (0.00) | (0.00) | (0.00) |
| Observations | 150,660 | 150,660 | 150,660 | 150,660 | 150,660 | 150,660 |
| R-squared | 0.48 | 0.49 | 0.49 | 0.40 | 0.40 | 0.40 |
| Day FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Ticker FE | Yes | Yes | Yes | Yes | Yes | Yes |

*Notes:* This table presents regression results for specifications 1 and 2. Columns 1 to 3 show results with attention as the dependent variable, columns 4 to 6 contain results with respect to sentiment as the dependent variable. All models are estimated using ticker and day fixed effects and with standard errors clustered on the ticker level. Standard errors are given in parentheses.

Table 6: Investor Attention, Sentiment, and Earnings Surprise

| | Attention | | | Sentiment | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Abs(Return) | 71.46** | | 65.84** | | | |
| | (28.82) | | (30.63) | | | |
| Abs(Surprise) | | 22.66 | -96.97 | | | |
| | | (51.07) | (160.60) | | | |
| Abs(Return) | | | 1,744.09 | | | |
| ×Abs(Surprise) | | | (2,129.59) | | | |
| Return | | | | 12.22*** | | 12.15*** |
| | | | | (3.67) | | (3.72) |
| Surprise | | | | | 11.93 | 0.38 |
| | | | | | (11.22) | (10.26) |
| Return×Surprise | | | | | | -393.69 |
| | | | | | | (324.82) |
| Constant | 4.89*** | 6.94*** | 5.13*** | -1.14*** | -1.15*** | -1.12*** |
| | (0.85) | (0.12) | (0.99) | (0.00) | (0.01) | (0.01) |
| Sample | EAD=1 | EAD=1 | EAD=1 | EAD=1 | EAD=1 | EAD=1 |
| Observations | 2,372 | 2,372 | 2,372 | 2,372 | 2,372 | 2,372 |
| R-squared | 0.81 | 0.81 | 0.81 | 0.75 | 0.74 | 0.75 |
| Day FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Ticker FE | Yes | Yes | Yes | Yes | Yes | Yes |

*Notes:* This table presents regression results for variations of specifications 1 and 2. The sample is restricted to observations with earnings announcements. Columns 1 to 3 show results with attention as the dependent variable, columns 4 to 6 contain results with respect to sentiment as the dependent variable. All models are estimated using ticker and day fixed effects and with standard errors clustered on the ticker level. Standard errors are given in parentheses.

Table 7: Cumulative Abnormal Returns Subsequent to Earnings Announcements

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | CAR[1,2] | CAR[1,3] | CAR[1,4] | CAR[1,2] | CAR[1,3] | CAR[1,4] |
| Surprise | 78.54*** | 106.46*** | 104.00*** | 87.72*** | 107.91*** | 98.43*** |
| | (21.67) | (32.79) | (32.91) | (20.98) | (34.36) | (31.07) |
| Sentiment | 0.04** | 0.03 | 0.01 | 0.12*** | 0.10** | 0.10** |
| | (0.02) | (0.03) | (0.03) | (0.04) | (0.03) | (0.05) |
| Surprise × Sentiment | -4.04 | -3.33 | -4.52 | -1.93 | -2.26 | -2.50 |
| | (2.46) | (4.70) | (5.15) | (2.42) | (6.41) | (6.68) |
| Constant | -0.07 | -0.07 | -0.18 | 0.01 | 0.01 | -0.09 |
| | (0.09) | (0.12) | (0.12) | (0.05) | (0.06) | (0.06) |
| Sample | EAD=1 | EAD=1 | EAD=1 | EAD=1 | EAD=1 | EAD=1 |
| Observations | 2,378 | 2,378 | 2,378 | 2,377 | 2,377 | 2,377 |
| R-squared | 0.12 | 0.11 | 0.11 | 0.32 | 0.30 | 0.29 |
| Day FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Ticker FE | No | No | No | Yes | Yes | Yes |

*Notes:* This table presents regression results for specification 3. Columns 1–3 report results with day fixed effects only, while columns 4–6 additionally include ticker fixed effects. Standard errors are clustered on the ticker level and are reported in parentheses.

Table 8: Cumulative Abnormal Returns for User Groups

| | (1) CAR[1,2] | (2) CAR[1,3] | (3) CAR[1,4] | (4) CAR[1,2] | (5) CAR[1,3] | (6) CAR[1,4] |
|---|---|---|---|---|---|---|
| Surprise | 78.18*** | 109.40*** | 105.28*** | 87.74*** | 111.96*** | 101.82*** |
| | (21.64) | (33.53) | (32.44) | (21.13) | (33.44) | (30.20) |
| Sent$\overline{\text{InfoContro}}$ | 0.19*** | 0.11 | 0.05 | 0.32*** | 0.22** | 0.15 |
| | (0.07) | (0.10) | (0.10) | (0.09) | (0.11) | (0.12) |
| Sent$\overline{\text{Info}}$Contro | 0.07 | 0.08 | 0.05 | 0.06 | 0.06 | 0.07 |
| | (0.07) | (0.11) | (0.13) | (0.06) | (0.12) | (0.14) |
| SentInfo$\overline{\text{Contro}}$ | 0.25 | 0.57* | 0.47 | 0.39** | 0.63** | 0.53 |
| | (0.18) | (0.29) | (0.33) | (0.19) | (0.32) | (0.36) |
| SentInfoContro | -0.10* | -0.19** | -0.16 | 0.00 | -0.07 | -0.04 |
| | (0.06) | (0.09) | (0.10) | (0.05) | (0.10) | (0.11) |
| Surprise×Sent$\overline{\text{InfoContro}}$ | -0.10 | 7.99 | 15.36 | 0.15 | 5.45 | 10.99 |
| | (13.35) | (24.07) | (23.40) | (12.24) | (25.28) | (25.05) |
| Surprise×Sent$\overline{\text{Info}}$Contro | -3.25 | 23.23 | 1.74 | 3.99 | 36.78 | 25.15 |
| | (31.80) | (73.53) | (66.80) | (27.76) | (77.38) | (76.57) |
| Surprise×SentInfo$\overline{\text{Contro}}$ | -73.06* | -103.91 | -91.32 | -99.81*** | -130.04** | -115.10* |
| | (39.31) | (64.98) | (62.54) | (32.38) | (64.45) | (66.51) |
| Surprise×SentInfoContro | 0.37 | -9.17 | -7.48 | 5.92 | -6.22 | -6.33 |
| | (11.36) | (19.31) | (18.57) | (9.89) | (21.20) | (21.59) |
| Constant | -0.07 | -0.06 | -0.18 | 0.01 | 0.01 | -0.09* |
| | (0.09) | (0.12) | (0.12) | (0.04) | (0.05) | (0.05) |
| Observations | 2,378 | 2,378 | 2,378 | 2,377 | 2,377 | 2,377 |
| R-squared | 0.12 | 0.12 | 0.12 | 0.33 | 0.30 | 0.30 |
| Day FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Ticker FE | No | No | No | Yes | Yes | Yes |

*Notes:* This table presents regression results from variants of specification 4, comparing models without ticker fixed effects (columns 1–3) and with both ticker and day fixed effects (columns 4–6). The sample is restricted to earnings announcement days. Robust standard errors clustered at the ticker level are reported in parentheses.

# Appendix

## A Prompt

You are an assistant that classifies Reddit comments from WallStreetBets as either **Informed** or **Uninformed**.

- **Informed** comments discuss accounting and finance topics, micro and macroeconomic developments, or regulations by the SEC.
- **Uninformed** comments do not discuss these topics.

**Examples of Informed Comments:**

1. "The company"s Q3 earnings report shows a significant increase in net income due to cost-cutting measures."
2. "Interest rate hikes by the Federal Reserve could negatively impact the housing market."
3. "New SEC regulations on cryptocurrency trading are likely to affect market liquidity."
4. "The merger between these two tech giants will create a monopoly, raising antitrust concerns."
5. "A weakening dollar could boost exports but may also lead to higher inflation."

**Examples of Uninformed Comments:**

1. "This stock is going to the moon! Buy now!"
2. "Just YOLO'd my life savings into options. Wish me luck!"
3. "Who's up for a trading challenge this week?"
4. "I don't care about fundamentals; I'm all about the hype."
5. "My friend's cousin said this stock will explode tomorrow."

**Task:**
Classify the following Reddit comment as either **Informed** or **Uninformed**.
**Reddit Comment:**
"{comment}"

Figure A.1: S&P500 Returns of Sample Firms

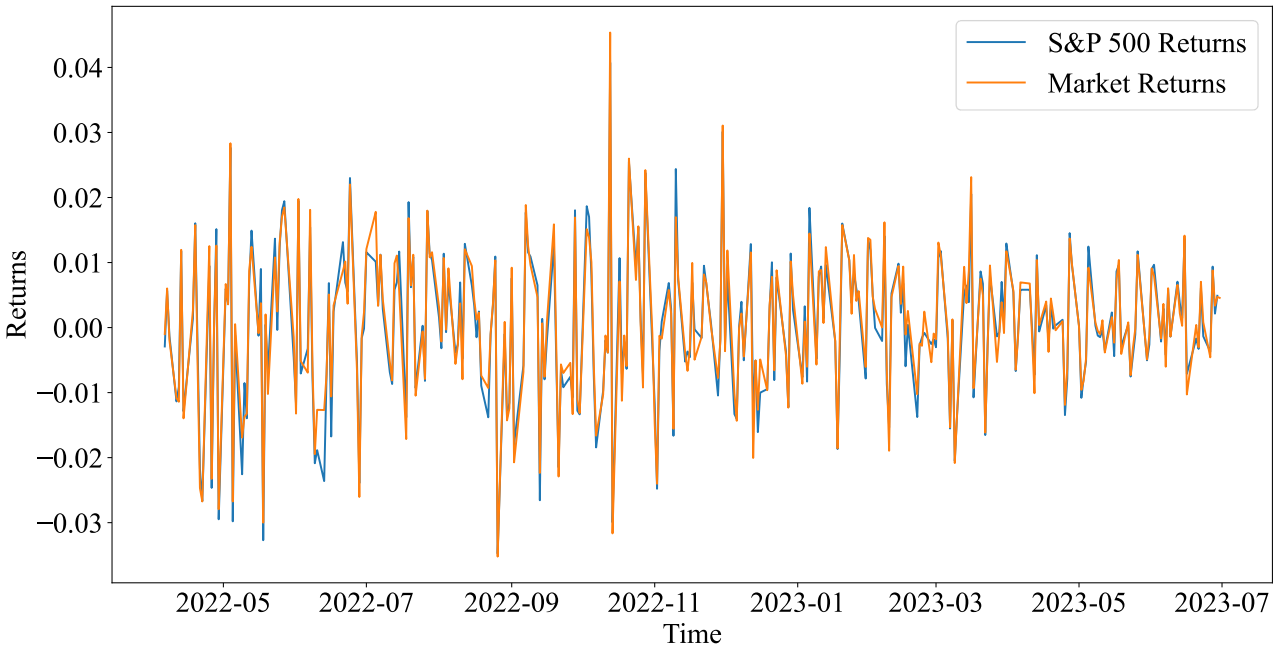


*Note:* Graphical display of the mean S&P 500 return (Yahoo Finance) and the return for all 486 firms in our sample over the whole period of 310 trading days. The orange line represents the S&P 500 return (Yahoo Finance), the blue line the return of firms included in the sample.

# B Tables

Table B.1: PPML Estimation of Attention on All Days and EAD Specifically

| | Attention – Full Sample | | | Attention – EAD = 1 | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| EAD | 1.70*** | | 1.72*** | | | |
| | (0.11) | | (0.14) | | | |
| Abs(Return) | | 14.28*** | 14.21*** | 4.70*** | | 4.23*** |
| | | (1.44) | (0.87) | (1.34) | | (1.51) |
| EAD×Abs(Return) | | | -3.07 | | | |
| | | | (1.98) | | | |
| Abs(Surprise) | | | | | 12.33*** | 6.55 |
| | | | | | (4.26) | (9.82) |
| Abs(Return) | | | | | | 41.19 |
| ×Abs(Surprise) | | | | | | (159.52) |
| Constant | 3.01*** | 2.73*** | 2.65*** | 4.27*** | 4.40*** | 4.27*** |
| | (0.01) | (0.04) | (0.03) | (0.05) | (0.01) | (0.06) |
| Sample | All | All | All | EAD=1 | EAD=1 | EAD=1 |
| Observations | 128,520 | 128,520 | 128,520 | 1,306 | 1,306 | 1,306 |
| Pseudo R-squared | 0.80 | 0.80 | 0.82 | 0.92 | 0.92 | 0.92 |
| Day FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Ticker FE | Yes | Yes | Yes | Yes | Yes | Yes |

*Notes:* This table presents regression results for a poisson pseudo maximum likelihood adaptation of specification 1, estimated with Correia et al. (2020). For columns 1 to 3 we use the full sample, the sample is restricted to observations with earnings announcements for the estimation of columns 4 to 6. In all columns the attention (count data) is the outcome. All models are estimated using ticker and day fixed effects and with standard errors clustered on the ticker level. Standard errors are given in parentheses. PPML relies upon partialling out the fixed effects, thus the number of observations used for the estimation decreases.

**TRR 266 Accounting for Transparency**

*Contact:*
Prof. Dr. Dr. h.c. Dr. h.c. Caren Sureth-Sloane
Paderborn University
Faculty of Business Administration and Economics
Department of Taxation, Accounting and Finance
Warburger Str. 100, 33098 Paderborn, Germany

trr266@mail.upb.de
www.accounting-for-transparency.de

Download ZEW Discussion Papers:

https://www.zew.de/en/publications/zew-discussion-papers

or see:

https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html
https://ideas.repec.org/s/zbw/zewdip.html

## IMPRINT

**ZEW – Leibniz-Zentrum für Europäische Wirtschaftsforschung GmbH Mannheim**
ZEW – Leibniz Centre for European Economic Research

L 7,1 · 68161 Mannheim · Germany
Phone +49 621 1235-01
info@zew.de · zew.de