Social Media Sentiment and Over-the-Counter Equity Returns

Masahiro Watanabe

June 30, 2025

Abstract

This study examines the relation between social media sentiment and the pricing of over-the-counter stocks. We measure the sentiment of messages on the largest social network for investors, Stocktwits, using large language models. The online platform and the trading venue share retail investors as their clientele. Stock prices hike with positive sentiment and reverse afterwards, while they plummet upon negative sentiment with little subsequent reversal. The explanatory power of sentiment quantified by natural language processing is not subsumed by the binary classification of sentiment by investors themselves. These results have implications for the informativeness of messages and investors' trading behavior.

1 Introduction

This study examines the connection between social media and over-the-counter (OTC) equity markets, both of which are characterized by the significant involvement of individual participants rather than institutions. Driven by technological innovation and commission-free trading platforms, the rise of retail investors has fundamentally altered the dynamics of markets that have been long dominated by institutional investors. With diverse motivations such as social trends and behavioral biases, retail investors can introduce distinctive patterns of volatility and price discovery. Understanding their impact is essential, as their collective actions can drive significant short-term price movements and increasingly influence broader market trends.

In recent decades, U.S. equity markets have seen a remarkable shift, marked by the rapid expansion of OTC trading. Although traditional trading venues like the New York Stock Exchange (NYSE) and the National Association of Securities Dealers Automatic Quotation System (NASDAQ) remain dominant in size (with market capitalizations of approximately $28 trillion and $25 trillion U.S. dollars, respectively), the OTC markets have grown swiftly, now surpassing $7 trillion dollars in capitalization and hosting thousands of domestic and international issuers. A key force behind this growth is the proliferation of social media platforms such as Stocktwits, which has democratized market information and amplified retail investor activity, especially in speculative OTC securities. Empowered by real-time news and collective sentiment, retail investors are increasingly shaping the OTC landscape, making it a critical, yet often overlooked, component of the U.S. equity ecosystem. This convergence of social media and retail trading raises important questions about market structure, investor protection, and regulation.

To gauge retail investor sentiment, we apply large language models to messages on Stocktwits, the largest social media specializing in investment. We employ simple pre-trained models based on the Bidirectional Encoder Representations from Transformers (BERT, Devlin et al. (2019)), which offer several advantages over generative artificial intelligence (AI) when it comes to sentiment analysis. BERT models are specifically designed for and excel at understanding the context and nuances in short texts, such as social media posts. They are faster to deploy, require less computational resources, and produce more consistent results for sentiment classification. In contrast, generative AI models are often more complex, resource-intensive, and may introduce prompt-dependent variability, making them less efficient for focused tasks like sentiment analysis. This makes simple BERT models a more practical and reliable choice for analyzing Stocktwits messages.

We analyze over 8 million Stocktwits posts from November 2020 to August 2023 and construct daily portfolios of OTC stocks. Prior to this period, Stocktwits largely banned the discussions of OTC stocks. Sentiment probabilities from three BERT models are updated in a Bayesian manner. Specifically, a BERT model gives categorical probabilities for sentiment labels. Starting with a diffuse Dirichlet prior, the posterior distribution of categorical probabilities given the data from BERT remains Dirichlet with updated concentration parameters, because the Dirichlet distribution is a conjugate prior of the categorical distribution. This allows us to update the distribution of sentiment as we process messages. Sentiment and neutrality measures are computed for each stock each day in this way. Thus, our framework offers a tractable probabilistic model of sentiment that evolves over time.

We observe that stock prices typically start rising before the posting of messages with positive sentiment. However, this increase is reversed within a month, suggesting that such positive messages are unlikely to contain value-relevant information. In contrast, stock prices experience a pronounced decline prior to the posting of messages with negative sentiment, with little to no reversal in the subsequent month. This pattern indicates that negative sentiment is more closely tied to adverse information about the firm's fundamental value. Given that OTC stocks are illiquid and difficult to sell short, these findings support the hypothesis that retail investors are incentivized to post positive messages about the stocks they own, even when these messages lack economic justification.

Building on this, we examine whether the strength of sentiment, as measured by text sharpness or neutrality, influences the effect of sentiment on returns. The evidence indicates that the effect of sentiment on returns is weaker when messages are textually neutral (dull), and stronger when the text is less neutral (sharper). This suggests that sharp, clearly expressed sentiment has a greater association on stock price movements.

We further address whether Stocktwits messages provide information not already captured by standard risk models. By adjusting stock returns for CAPM risk, the analysis shows that the reversal after positive sentiment posts persists even after accounting for risk, implying that these price moves are not explained by risk factors and may instead reflect mispricing. In contrast, negative sentiment stocks do not exhibit significant reversals.

Our regression analyses control for a variety of characteristics, including investors' own labeling of sentiment as bullish or bearish. Thus, the explanatory power of sentiment quantified by natural language processing is not captured by this binary classification or other stock characteristics.

The rest of the paper is organized as follows. The next section provides a brief literature review. Section 3 explains the data and methodology. Section 4 presents the results. Section 5 concludes.

2 Literature Review

The past decade has witnessed an unprecedented surge in retail trading activity, fueled in part by the democratization of information through social media platforms such as X (formerly Twitter), Reddit, and StockTwits. Social media has fundamentally altered the information environment, enabling retail investors to coordinate, share ideas, and react swiftly to news and rumors. Researchers have begun to document the tangible effects of social media activity on market outcomes, particularly for less liquid or more opaque securities, such as those traded over the counter.

Barber, X. Huang, and Odean (2022) provide robust evidence on the ability of retail investors to move prices, especially in stocks with limited institutional coverage or low liquidity. Their findings suggest that social media-fueled herding can lead to significant price dislocations, particularly in smaller-capitalization stocks. They find that intense buying by Robinhood users forecasts negative returns, which is consistent with the reversal of price appreciation we document for stocks with positive sentiment. Similarly, Bryzgalova, Pavlova, and Sikorskaya (2023) analyze the "meme stock" phenomenon, where retail traders—often galvanized by viral social media posts—collectively drive up prices of certain stocks via option trading, sometimes detached from fundamental valuations.

Cookson, Engelberg, and Mullins (2022) report evidence of "echo chambers" on Stocktwits, where individuals selectively expose themselves to information that is confirmatory to their self-declared sentiment labels. Beliefs formed in echo chambers are associated with larger trading volume and lower ex-post returns. Cookson and Niessner (2020) document that investor disagreement, measured by social media sentiment and self-classified investment styles, stems equally from differences in information and interpretation. However, information-based disagreement is more strongly linked to trading volume, indicating it drives trading activity more than differences in market approach. Jiao, Veiga, and Walther (2020) find that coverage by traditional news media predicts decreases in subsequent stock price volatility and turnover, while coverage by social media forecasts increases in volatility and turnover.

None of these studies employs large language models to explore the impact of social media sentiment on stock prices. This is what we now turn to.

3 Data and Methodology

3.1 Data

We download 8,310,540 messages from Stocktwits, ten of which are excluded from the analysis due to empty text. Daily data on individual OTC stocks are provided by the OTC Markets Group. We extract active U.S. common stocks on non-Grey Market with Class A or a missing class and no caveat emptor. Similarly to Eraker and Ready (2015), we require a minimum monthly volume of \$2,000, a market capitalization of \$1 million, and a minimum price of \$0.01 within a month. The two datasets are merged by ticker symbol and date. The final portfolio returns consist of daily observations from November 3, 2020 to August 4, 2023. The beginning date of the sample is dictated by the fact that Stocktwits largely banned the discussions of OTC stocks prior to that day.

3.2 Bayesian Updating of Sentiment Probabilities

We employ three large language models based on BERT: FinBERT (Araci 2019), FinBERT-tone (A. H. Huang, Wang, and Yang 2022), and Sentiment-xDistil (Holmen 2023). FinBERT is trained on the Financial PhraseBank by Malo et al. (2014), while FinBERT-tone is trained on 10-K and 10-Q corporate reports, earnings call transcripts, and analyst reports. Sentiment xDistil is trained on news headlines and tweets annotated by ChatGPT 3.5. These models are pre-trained, and we do not train them further. This should help reproduce our result.

Each Stocktwits message has unobservable sentiment, denoted by a random variable, s. A BERT model gives the probabilities of a message's sentiment over three categories, "Negative," "Neutral," and "Positive," denoted by $(-, \circ, +)$, respectively. Figure 1 shows the ternary plot of these categorical probabilities from FinBERT. Most messages are concentrated around the Neutral category, implying that they are uninformative. However, there are significant masses near the Negative and Positive categories. We will see that sentiment computed from these probabilities are strongly associated with contemporaneous returns.

Let the vector of categorical probabilities be $\boldsymbol{p} = (p_-, p_\circ, p_+)$. Then s has a categorical distribution with probability vector $\boldsymbol{p}$. In a Bayesian framework, $\boldsymbol{p}$ also has a probability distribution, whose density is updated as we observe data. Since the probabilities add up to one,

$$p_{-} + p_{\circ} + p_{+} = 1, \tag{1}$$

their multivariate density can be modeled by what is known as the Dirichlet distribution. We denote its concentration parameter by $\boldsymbol{\alpha} = (\alpha_{-}, \alpha_{\circ}, \alpha_{+})$.

Start with a diffuse Dirichlet prior, $\boldsymbol{\alpha}_0 = (1, 1, 1)$. The next section illustrates this density. Draw a Stocktwits message and run a BERT model to get the first data, $s_1 \sim Cat\,(3, \boldsymbol{p}_1)$, which denotes the categorical distribution with three categories and the categorical probability vector $\boldsymbol{p}_1 = (p_{1-}, p_{1\circ}, p_{1+})$. Since the Dirichlet distribution is the conjugate prior of the categorical distribution, the posterior distribution of the categorical probabilities, $\boldsymbol{p}$, given the data, s_1, is also Dirichlet with concentration parameter $\boldsymbol{\alpha}_0 + \boldsymbol{p}_1$:

$$\begin{aligned} \boldsymbol{p} \mid s_1 &\sim \mathrm{Dir}\,(3, \boldsymbol{\alpha}_0 + \boldsymbol{p}_1) \\ &= \mathrm{Dir}\,(3, 1 + p_{1-}, 1 + p_{1\circ}, 1 + p_{1+}) \end{aligned} \tag{2}$$

where $\mathrm{Dir}(3, \cdot)$ denotes the Dirichlet distribution with three categories and the concentration parameters in the remaining argument. Repeat this over $J \geq 1$ messages. The posterior is now

$$\begin{aligned} \boldsymbol{p} \mid \{s_j\}_{j=1,\ldots,J} &\sim \mathrm{Dir}\left(3, \boldsymbol{\alpha}_0 + \sum_{j=1}^{J} \boldsymbol{p}_j\right) \\ &= \mathrm{Dir}\left(3, 1 + \sum_{j=1}^{J} p_{j-}, 1 + \sum_{j=1}^{J} p_{j\circ}, 1 + \sum_{j=1}^{J} p_{j+}\right) \end{aligned} \tag{3}$$

Equation (3) can also be interpreted as the posterior distribution after observing $\sum_{j=1}^{J} p_{j-}$, $\sum_{j=1}^{J} p_{j\circ}$, and $\sum_{j=1}^{J} p_{j+}$ pseudocounts of the Negative, Neutral, and Positive categories, respectively, if sentiment were observable. This interpretation is valid because the Dirichlet distribution is also the conjugate prior of the multinomial distribution.

3.3 Illustration

Figure 2 plots the prior and posterior densities of the Dirichlet distribution. Panel (A) is the diffuse prior with concentration parameter $\boldsymbol{\alpha}_0 = (1, 1, 1)$. The density is flat over the triangle defined by $p_{0-} + p_{0+} \leq 1, p_{0-} \geq 0, p_{0+} \geq 0$, reflecting the lack of prior knowledge. Given a point on this triangle, the neutral probability can be inferred by $p_{0\circ} = 1 - p_{0-} - p_{0+}$.

Panel (B) shows the posterior density for ticker symbol VENG on January 23, 2023, where the concentration parameter is computed from Equation (3) assuming that the day starts with a diffuse prior. 25 messages about this stock were posted on Stocktwits on that day. The posterior belief inferred from those messages is negatively skewed as visible in the figure. Note that the stock recorded the lowest daily excess return, -92%, in our sample on this day. Similarly, Panel (C) depicts a positively skewed posterior density for ticker symbol NNMX on January 4, 2022. The stock earned the highest excess return, 731%, in our sample and received 35 messages on Stocktwits on that day.

These illustrative examples suggest an association between stock return and BERT sentiment probabilities. The following sections will demonstrate the strength of this association.

3.4 Computing Sentiment

Attach the values, $-1, 0$, and 1, to the Negative, Neutral, and Positive categories, respectively. Then sentiment s takes these values with probabilities p_-, $p_\circ$, and p_+, respectively. Therefore, the expected sentiment of the j 'th message about stock i on day t is

$$\mathbb{E}\left[s_{j,i,t}\right] = p_{j,i,t+} - p_{j,i,t-}, \tag{4}$$

We define the sentiment measure, $SENT_{i,t}$, as the average expected sentiment of $J_{i,t}$ messages about stock i on day t and the text neutrality measure, $NEUT_{i,t}$, as the average probability of the Neutral category:

$$SENT_{i,t} = \frac{1}{J_{i,t}} \sum_{j=1}^{J_{i,t}} \mathbb{E}\left[s_{j,i,t}\right] = \frac{1}{J_{i,t}} \sum_{j=1}^{J_{i,t}} \left(p_{j,i,t+} - p_{j,i,t-}\right) \tag{5}$$

$$NEUT_{i,t} = \frac{1}{J_{i,t}} \sum_{j=1}^{J_{i,t}} p_{j,i,t\circ} \tag{6}$$

These variables assume values in the ranges, $-1 \leq SENT_{i,t} \leq 1$ and $0 \leq NEUT_{i,t} \leq 1$.

4 Results

4.1 Sentiment, Neutrality, and Stock Returns

Figure 3 plots the histograms of the sentiment ($SENT$) and neutrality ($NEUT$) measures as well as the probabilities of the three labels as classified by each of the three BERT models. Panel A shows that $SENT$ from FinBERT has a sharp mass around a slightly positive mode. According to Panel A of Table 1, mean sentiment is 0.0318 with a gigantic t-statistic of 302.4, implying that Stocktwits messages are slightly positive on average. The sharpness of the distribution comes from that of the negative and positive probabilities: their masses are concentrated around zero, implying that most messages are neither strongly negative nor strongly positive as seen already in Figure 1. However, they are consistently slightly positive, leading to the huge t-statistic. This is mirrored in the concentration of probabilities for the neutral label around one. Note that the neutral label does not contribute to the value of SENT , because it is assigned the value of zero.

Panels B and C show the corresponding histograms for FinBERT-tone and Sentiment xDistil. The distributions are even sharper than FinBERT; note that the y-axis scale is an order of magnitude larger than Panel A. This suggests the need for care in their use. Nevertheless, we will see that the sentiment and neutrality measures give qualitatively similar pricing results except where noted.

Table 2 tabulates the mean excess return by the sign of sentiment. For all the three BERT models, the mean excess return is negative for negative sentiment stocks ($SENT < 0$, $S = 1$) and positive for positive sentiment stocks ($SENT > 0$, $S = 2$). The means are statistically significant with the absolute t-statistics 9 or larger except for $S = 1$ for FinBERT-tone, which still has a significant t-statistic of -2.6.

4.2 Stock Returns around Stocktwits Posts

To examine the pricing of Stocktwits posts, we sort stocks by the sentiment ($SENT$) and neutrality ($NEUT$) measures independently every day (denoted as Day 0). We form two portfolios based on the sign of $SENT$, one with negative sentiment ($SENT < 0$, denoted as

S1) and the other with positive sentiment ($SENT > 0$, S2). Separately, we also form four portfolios as the intersection of the sign of $SENT$ and the $NEUT$ being below (denoted as $N1$) or above ($N2$) the median. Figure 4 plots the cumulative excess returns of these equally weighted portfolios from calendar days -10 through 30 using $SENT$ and $NEUT$ from the three BERT models. Tables 3 and 4 report the equally weighted mean excess return of constituent stocks within the time window shown in the rows starting with "Days", with the corresponding t-statistic beneath it.

The plots in Panel A of Figure 4 exhibit pre-event drifts: starting ten days prior to Stocktwits-posting, the cumulative return hikes for the positive sentiment portfolio (in orange) and plunges for the negative sentiment portfolio (in blue) toward the posting day, although the magnitude differs across the three BERT models.[1] According to Table 3, the mean returns for the positive sentiment portfolio over Days -10 to 0 for FinBERT and Sentiment-xDistil are 0.264% (t = 11.8) and 0.223% per day (t = 10.4), respectively. These are not only statistically significant, but also economically large: the annualized returns will be 94.3% and 75.3% per year, respectively. In contrast, the mean return is insignificant for FinBERT-tone. The negative sentiment portfolio experiences an even more dramatic price movement: the mean return over Days -10 to 0 in Table 3 is -0.848% per day (t = -31.6) for FinBERT, -0.678% (t = -20.9) for FinBERT-tone, and -0.969% (t = -33.7) for Sentiment-xDistil, compounding to gigantic annualized returns of -88.3%, -82.0%, and -91.4% per year, respectively.

Interestingly, the cumulative return on the positive sentiment portfolio reverses after the

[1] The somewhat rough plots appear to exhibit a weekly calendar effect.

posting day regardless of the model to measure sentiment. For example, the mean returns over Days 1 to 7, 8 to 14, 15 to 21, and 22 to 28 are -0.19% (t = -7.0), -0.16% (t = -6.0), -0.15% (t = -5.6), and -0.10% (t = -3.8), respectively, for FinBERT, and similarly negative and significant for the other two language models. The prolonged negative returns result in a reversal and even lead to an overshoot for longer horizons, as depicted in Figure 4. Therefore, the contents of positive Stocktwits posts are unlikely to be informative.

In contrast, the price of the negative sentiment portfolio remains depressed for FinBERT and Sentiment-xDistil. This is visible in Panel A of Figure 4, where the plots in blue are mostly flat after Day 0 for these two models. Consistently, the mean returns for Portfolio S1 after the posting day are insignificant in Table 3. Although the plot for the negative sentiment portfolio with FinBERT-tone weakly reverses after the posting day, Table 3 implies that this reversal is only marginally significant, as indicated by modest t-statistics of around 1.9 for the mean return on Portfolio S1 over Days 8 to 14 and 15 to 21. This is due to relatively large standard errors owing to the scarcity of negative sentiment for this language model.[2]

4.3 Text Sharpness Strengthens Sentiment

Panel B of Figure 4 plots the cumulative returns of sentiment-neutrality portfolios. It indicates that the effect of sentiment is weaker for more neutral text. Equivalently, text sharpness, defined as the complement of neutrality, strengthens sentiment: the surge for positive sentiment stocks and the dip for negative sentiment stocks prior to Stocktwits posting are

[2]For instance, although the standard deviation of excess returns over Days 1 to 7 is similar between the negative and positive sentiment portfolios at 0.103 and 0.095 for FinBERT-tone (unreported in a table), the number of observations is 48,026 and 146,409, respectively. These numbers together give the standard errors of 0.0468% and 0.0248%. With the mean returns of -0.031% and -0.149% shown in Table 3, they lead to the t-statistics of -0.66 and -6.0 there.

larger with sharp text ($NEUT$ is below the median, plotted in green and blue) than with dull text ($NEUT$ is above the median, plotted in red and orange) for all the three BERT models. In fact, there is no discernible run-up for the portfolio with positive sentiment and neutral text in red. Formally, the mean excess return for the S2N2 stocks over Days -10 to 0 in Table 4 is insignificant for FinBERT and Sentiment-xDistil, and is puzzlingly negative and signficant for FinBERT-tone; we would expect it to be positive, if any. This suggests a problem in applying the model trained on long text to short messages; recall that FinBERT-tone is trained on 10-K, 10-Q, earnings calls, and analyst reports. Similarly, the plot for the sharp negative sentiment portfolio in blue plummets more than the dull counterpart in orange prior to the day of posting for all the three models. The negative mean return on S1N1 stocks during Days -10 to 0 in Table 4 is negative and larger in magnitude by about 50% or more than S1N2 stocks across the three BERT models.

4.4 Are Messages Informative?

The previous section suggests that positive Stocktwits messages are unlikely to be infor mative, because the pre-event drift reverses. How about negative messages? To examine this question, on each day t we regress each individual stock's return, R_i, on the over-the-counter market return, R_M, using observations between days $t-60$ to $t-11$ requiring at least 30 observations:

$$R_{i,s} = \alpha_{i,t} + \beta_{i,t} R_{M,s} + \varepsilon_{i,s}, \quad t-60 \leq s \leq t-11. \tag{7}$$

Then using the estimated coefficients $\alpha_{i,t}$ and $\beta_{i,t}$, we adjust the individual stock return for the CAPM risk by subtracting the fitted return from the return over days $t-10$ through $t+30$,

$$R^a_{i,s} = R_{i,s} - \alpha_{i,t} - \beta_{i,t} R_{M,s}, \quad t-10 \leq s \leq t+30. \tag{8}$$

Table 5 shows the mean risk-adjusted (abnormal) return in Equation (8) separately on negative (S1) and positive (S2) sentiment stocks over the periods shown in the first column. The properties of returns prior to Stocktwits-posting remain similar to Table 3 after the risk adjustment: the risk-adjusted mean return is significantly negative on negative sentiment stocks across the three BERT models and significantly positive on positive sentiment stocks except for FinBERT-tone. Likewise, the risk-adjusted mean return on positive sentiment stocks after the posting of messages are significantly negative regardless of the BERT model.

Panel A of Figure 5 plots the cumulative abnormal return (CAR) of the two sentiment portfolios. Since the OTC market return was positive on average during our sample period, with a positive loading on the market, the CAR of positive sentiment portfolio consistently drifts downward. The CAR of negative sentiment portfolio, however, is similar to the risk-unadjusted plot in Figure 4: if anything, the plot is now smoother due to risk-adjustment. This is because the comovement between the individual stocks and the aggregate market cancels out in Equation (8) and makes the resulting $R^a_{i,s}$ less volatile.

Table 6 summarizes the mean risk-adjusted return on stocks in the intersections of sentiment and neutrality groups. Like the mean excess returns in Table 4, the mean risk-adjusted returns before Day 0 are significant and carry the same sign as sentiment except for the positive sentiment stocks with dull text (S2N2) and the sharp, positive-sentiment stocks (S2N1 in Panel B) as classified by FinBERT-tone. The positive-sentiment stocks with a significant run-up prior to Stocktwits posting (S2N1 in Panels A and C) continue to exhibit reversal afterwards upon risk adjustment. Reversals on negative sentiment stocks tend to be weaker.

Figure 6 shows the bi-grams of aggregate Stocktwits posts, and Figures 7 and 8 their breakdown by the messages' sentiment-neutrality levels and by the self-declared expertise of the users who post the messages, respectively. Tables A1, A2, and A3 in the appendix tabulate the bigrams by frequency. Focusing on FinBERT in Panel A of Figure 7, some of the words associated with negative sentiment in the top row, such as "delayed filed" and "filed SEC", express potentially serious concerns. Figure 8 reveals that these are typically posted by professionals rather than novice or intermediate users. In contrast, words with positive sentiment in the bottom row of Figures 7 are mixed in tone, as exemplified by the equally frequent occurrences of "close decreased" and "close increased". This is consistent with the previous observation that low returns on negative sentiment stick while high returns on positive sentiment reverse within a month.

Overall, the evidence in this section suggests that the response of OTC stocks to Stocktwits posts cannot be explained by the CAPM risk. It is more likely due to mispricing than other sources of risk for stocks with positive sentiment, because the pre-event drift com pletely reverses within several weeks. In other words, positive messages are uninformative. Text sharpness enhances sentiment.

4.5 Panel Regressions

Table 7 reports the result of panel regressions controlling for various stock characteristics beyond the CAPM risk. The dependent variable is the excess return on individual stocks. The sentiment measure has an extremely significant, positive coefficient across all specifications over the three BERT models. For example, the SENT coefficient in the full model (4) of Panels A, B, and C is 0.0655 (t = 15.4), 0.0289 (t = 10.5), and 0.0500 (t = 15.8), respectively, where the t-statistics in the parentheses above are computed using the standard errors in the table (with slightly more precision). This implies that a one-percent increase in sentiment is associated with a 2.9 to 6.6 basis-point increase in excess return. For the FinBERT and Sentiment-xDistil models, the coefficients on the four lagged sentiment measures, $SENT_{m1}$, $SENT_{m5-m2}$, $SENT_{m10-m6}$, and $SENT_{m30-m11}$, are significantly negative and interestingly, tend to become larger in magnitude as the lag length increases. Specifically, the coefficient estimates on those four lags for FinBERT are -0.0106 (t = -3.1), -0.0100 (t = -2.2), -0.0194 (t = -3.7), and -0.0249 (t = -2.6), respectively, where the t-statistics in the parentheses are computed similarly. They sum up to -0.065, which almost equals the coefficient on contemporaneous $SENT$, 0.0655. This implies that a return associated with a shock in sentiment on a given day will be fully reversed within 30 days if a shock of the same sign and magnitude occurs every day, consistent with our previous finding using sentiment portfolios. Considering the large spread on OTC stocks, therefore, investors who trade in the same direction as the tone of Stocktwits posts are unlikely to make significant profits on average.

5 Conclusion

This paper explores the intersection between social media sentiment and the pricing of over-the-counter (OTC) equities, a domain where retail investors are especially active and information asymmetries are pronounced. Using over 8 million Stocktwits posts, we apply domain-specific BERT models to quantify the sentiment and neutrality of messages, and then link these measures to OTC stock returns from November 2020 to August 2023.

Our main findings can be summarized as follows. First, we document that positive sentiment posts on Stocktwits are typically preceded by rising OTC stock prices, but this run-up tends to reverse over the following month. This pattern is robust across multiple BERT-based sentiment models and persists after adjusting for CAPM risk and stock characteristics. The reversal is concentrated in stocks with sharp, unambiguous sentiment; in contrast, neutral or dull messages exhibit weaker price effects. Notably, the reversal is much less apparent for negative sentiment posts, suggesting an asymmetry in how the market responds to different types of sentiment.

Second, we show that the sentiment extracted from large language models is not subsumed by investors' self-labeled sentiment tags. Panel regressions confirm that machine-quantified sentiment is an economically and statistically significant predictor of short-term returns of positive sentiment stocks, even after controlling for a range of observable characteristics.

Third, by modeling message sentiment as a Bayesian updating process with a Dirichlet prior, we provide a principled framework for aggregating probabilistic model outputs and for quantifying the uncertainty inherent in text-based sentiment analysis. This approach improves interpretability and can be extended to other settings with multi-class outputs.

Taken together, our evidence suggests that social media sentiment, as measured by modern large language models, is associated with temporary mispricing in OTC stocks, likely reflecting retail-driven excess optimism that corrects over time. This highlights the role of social platforms in both propagating information and amplifying behavioral biases in less regulated, less liquid markets.

Several promising avenues for future research emerge from our work. First, while we document strong associations between sentiment and returns, establishing causality remains a challenge. Future studies could use exogenous shocks to social media activity (such as platform outages or regulatory interventions) to better identify the causal impact of sentiment on OTC pricing.

Second, it is unclear whether sentiment-driven effects observed in OTC stocks spill over to listed markets, or vice versa. Studies could examine cross-market linkages and contagion channels, especially during episodes of heightened retail activity.

Third, given the potential for rapid sentiment shifts to induce mispricing, there is scope for developing real-time monitoring tools for regulators and exchanges. Research could evaluate the effectiveness of such tools in mitigating manipulation or promoting market integrity.

Finally, while our study uses classification-based BERT models, generative AI models are rapidly advancing. Future work could compare the predictive power and interpretability of generative versus discriminative models for sentiment analysis in financial contexts.

By addressing these open questions, the literature can better elucidate the ways in which social media—and the NLP models used to parse it—reshape the informational efficiency, volatility, and fairness of modern financial markets.

References

Araci, Dogu (2019). "FinBERT: Financial Sentiment Analysis with Pre-trained Language Models". In: *arXiv, https://arxiv.org/abs/1908.10063, model downloaded from https://huggingface.co/ProsusAI/finbert*.

Barber, Brad, Xing Huang, and Terrance Odean (2022). "Attention Induced Trading and Returns: Evidence from Robinhood Users". In: *Journal of Finance* 77.6, pp. 3141–3190.

Bryzgalova, Svetlana, Anna Pavlova, and Taisiya Sikorskaya (2023). "Retail Trading in Options and the Rise of the Big Three Wholesalers". In: *Journal of Finance* 78.6, pp. 3465–3514.

Cookson, J. Anthony, Joseph E Engelberg, and William Mullins (Aug. 2022). "Echo Chambers". In: *The Review of Financial Studies* 36.2, pp. 450–500. ISSN: 0893-9454. DOI: 10.1093/rfs/hhac058. eprint: https://academic.oup.com/rfs/article-pdf/36/2/450/48744228/hhac058.pdf. URL: https://doi.org/10.1093/rfs/hhac058.

Cookson, J. Anthony and Marina Niessner (2020). "Why Don't We Agree? Evidence from a Social Network of Investors". In: *The Journal of Finance* 75.1, pp. 173–228. DOI: https://doi.org/10.1111/jofi.12852. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12852. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12852.

Devlin, Jacob et al. (2019). "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding". In: *arXiv, https://arxiv.org/abs/1810.04805*.

Eraker, Bjorn and Mark Ready (2015). "Do investors overpay for stocks with lottery-like payoffs? An examination of the returns of OTC stocks". In: *Journal of Financial Economics*

115.3, pp. 486–504. ISSN: 0304-405X. DOI: https://doi.org/10.1016/j.jfineco.2014.11.002. URL: https://www.sciencedirect.com/science/article/pii/S0304405X14002372.

Holmen, Hakon Magne (2023). "Sentiment-xDistil". In: *Hugging Face, model downloaded from https://huggingface.co/hakonmh/sentiment-xdistil-uncased*.

Huang, Allen H., Hui Wang, and Yi Yang (2022). "FinBERT: A Large Language Model for Extracting Information from Financial Text". In: *Contemporary Accounting Research, model downloaded from https://huggingface.co/yiyanghkust/finbert-tone*.

Jiao, Peiran, Andre Veiga, and Ansgar Walther (2020). "Social media, news media and the stock market". In: *Journal of Economic Behavior and Organization* 176, pp. 63–90. ISSN: 0167-2681. DOI: https://doi.org/10.1016/j.jebo.2020.03.002. URL: https://www.sciencedirect.com/science/article/pii/S0167268120300676.

Malo, Pekka et al. (Apr. 2014). "Good Debt or Bad Debt: Detecting Semantic Orientations in Economic Texts". In: *Journal of the American Society for Information Science and Technology*. DOI: 10.1002/asi.23062.

Table 1: Summary Statistics

Panel A: FinBERT

	Negative	Neutral	Positive	SENT
count	8,310,530	8,310,530	8,310,530	8,310,530
mean	0.111	0.747	0.143	0.032
(t-stat)	(1,641.3)	(8,592.3)	(2,067.1)	(302.4)
std	0.195	0.250	0.199	0.304
min	0.006	0.008	0.006	-0.970
25%	0.020	0.696	0.040	-0.024
50%	0.034	0.863	0.065	0.027
75%	0.080	0.912	0.134	0.098
max	0.977	0.960	0.962	0.948

Panel B: FinBERT-tone

	Negative	Neutral	Positive	SENT
count	8,310,530	8,310,530	8,310,530	8,310,530
mean	0.038	0.824	0.138	0.099
(t-stat)	(661.8)	(6,702.1)	(1,215.1)	(757.3)
std	0.167	0.355	0.326	0.378
min	0.000	0.000	0.000	-1.000
25%	0.000	0.967	0.000	-0.000
50%	0.000	1.000	0.000	-0.000
75%	0.001	1.000	0.002	0.000
max	1.000	1.000	1.000	1.000

Panel C: Sentiment-xDistil

	Negative	Neutral	Positive	SENT
count	8,310,530	8,310,530	8,310,530	8,310,530
mean	0.096	0.733	0.171	0.075
(t-stat)	(1,051.7)	(5,286.7)	(1,439.0)	(466.0)
std	0.263	0.400	0.342	0.461
min	0.000	0.000	0.000	-1.000
25%	0.000	0.410	0.000	-0.000
50%	0.001	0.993	0.001	0.000
75%	0.006	1.000	0.048	0.024
max	1.000	1.000	1.000	1.000

This table shows the summary statistics for the three BERT models shown in the panel titles. "Negative," "Neutral," and "Positive" are the probabilities assigned to those labels. $SENT$ and AMB are the sentiment and neutrality measures in Equations (5) and (6), respectively.

Table 2: Excess Return by Sentiment

Panel A: FinBERT

	S	exRET	SENT	NEUT
count	1	45,344	46,436	46,436
mean	1	-0.010	-0.128	0.740
(t-stat)	1	(-17.3)	(-148.4)	(827.1)
count	2	84,527	85,875	85,875
mean	2	0.017	0.136	0.761
(t-stat)	2	(35.1)	(223.2)	(1,215.7)

Panel B: FinBERT-tone

	S	exRET	SENT	NEUT
count	1	42,668	43,658	43,658
mean	1	-0.002	-0.051	0.932
(t-stat)	1	(-2.6)	(-70.4)	(1,151.2)
count	2	87,203	88,653	88,653
mean	2	0.013	0.178	0.789
(t-stat)	2	(26.7)	(210.1)	(922.3)

Panel C: Sentiment-xDistil

	S	exRET	SENT	NEUT
count	1	39,999	40,983	40,983
mean	1	-0.006	-0.263	0.652
(t-stat)	1	(-9.6)	(-159.6)	(399.1)
count	2	89,872	91,328	91,328
mean	2	0.014	0.214	0.725
(t-stat)	2	(31.0)	(234.0)	(773.4)

This table shows the summary statistics by sentiment for the three BERT models shown in the panel titles. $S = 1$ is the stocks with negative sentiment ($SENT \leq 0$) and $S = 2$ those with positive sentiment ($SENT > 0$). $exRET$ is the excess return. $SENT$ and $NEUT$ are the sentiment and neutrality measures in Equations (5) and (6), respectively.

Table 3: Excess Returns on Sentiment Portfolios around Stocktwits Post

Panel A: FinBERT

	S1	S2
Days -10 to 0	-0.00849	0.00265
t	-31.61703	11.78248
Days 1 to 7	0.00000	-0.00190
t	0.01211	-7.02373
Days 8 to 14	-0.00022	-0.00163
t	-0.57685	-6.02082
Days 15 to 21	-0.00054	-0.00151
t	-1.40841	-5.64476
Days 22 to 28	-0.00044	-0.00103
t	-1.13912	-3.80015

Panel B: FinBERT-tone

	S1	S2
Days -10 to 0	-0.00678	0.00024
t	-20.86740	1.15725
Days 1 to 7	-0.00031	-0.00148
t	-0.65521	-5.98127
Days 8 to 14	0.00099	-0.00179
t	1.95332	-7.28724
Days 15 to 21	0.00096	-0.00184
t	1.97445	-7.49580
Days 22 to 28	0.00024	-0.00115
t	0.50128	-4.61564

Panel C: Sentiment-xDistil

	S1	S2
Days -10 to 0	-0.00969	0.00223
t	-33.72450	10.38117
Days 1 to 7	-0.00009	-0.00170
t	-0.20422	-6.59725
Days 8 to 14	-0.00021	-0.00151
t	-0.48101	-5.82721
Days 15 to 21	-0.00013	-0.00161
t	-0.30161	-6.23546
Days 22 to 28	-0.00013	-0.00111
t	-0.31458	-4.29491

Stocks are sorted by the sentiment measure ($SENT$) into portfolios every day (Day 0) between November 2, 2020 and August 4, 2023. S1 is the negative sentiment portfolio ($SENT \leq 0$) and S2 the positive sentiment portfolio ($SENT > 0$). The rows starting with "Days" report the equally weighted mean excess return of constituent stocks within the time window indicated, with the corresponding t-statistic beneath it. Each panel uses $SENT$ from the BERT model shown in the title.

Table 4: Excess Returns on Sentiment-Neutrality Portfolios around Stocktwits Post

Panel A: FinBERT

	S1N1	S1N2	S2N1	S2N2
Days -10 to 0	-0.00941	-0.00576	0.00447	-0.00049
t	-30.48857	-10.60373	15.14193	-1.45036
Days 1 to 7	0.00006	-0.00017	-0.00211	-0.00154
t	0.15099	-0.21723	-6.26077	-3.38552
Days 8 to 14	-0.00043	0.00037	-0.00211	-0.00077
t	-0.94358	0.47481	-6.34556	-1.68335
Days 15 to 21	-0.00075	0.00009	-0.00205	-0.00056
t	-1.72086	0.11179	-6.27417	-1.21780
Days 22 to 28	-0.00070	0.00034	-0.00128	-0.00058
t	-1.58319	0.43697	-3.82186	-1.28058

Panel B: FinBERT-tone

	S1N1	S1N2	S2N1	S2N2
Days -10 to 0	-0.00914	-0.00489	0.00058	-0.00154
t	-18.33483	-11.43180	2.58221	-3.17753
Days 1 to 7	-0.00082	0.00010	-0.00161	-0.00081
t	-1.18526	0.16057	-6.00332	-1.26108
Days 8 to 14	0.00058	0.00132	-0.00214	0.00004
t	0.72858	2.02279	-8.05784	0.06655
Days 15 to 21	0.00119	0.00077	-0.00213	-0.00033
t	1.58141	1.21968	-8.11406	-0.49232
Days 22 to 28	0.00016	0.00031	-0.00119	-0.00090
t	0.21726	0.48263	-4.41146	-1.44010

Panel C: Sentiment-xDistil

	S1N1	S1N2	S2N1	S2N2
Days -10 to 0	-0.01178	-0.00702	0.00371	0.00025
t	-30.13156	-16.62599	12.39571	0.83198
Days 1 to 7	0.00049	-0.00082	-0.00189	-0.00144
t	0.92328	-1.23754	-5.51881	-3.69522
Days 8 to 14	0.00034	-0.00089	-0.00182	-0.00110
t	0.57043	-1.45683	-5.34795	-2.74125
Days 15 to 21	-0.00022	-0.00000	-0.00198	-0.00112
t	-0.40761	-0.00194	-5.97057	-2.73321
Days 22 to 28	-0.00017	-0.00008	-0.00146	-0.00064
t	-0.30995	-0.12706	-4.30241	-1.60958

Stocks are sorted independently by the sentiment ($SENT$) and neutrality ($NEUT$) measures into portfolios every day (Day 0) between November 2, 2020 and August 4, 2023. Four portfolios are formed as the intersection of negative ($SENT \leq 0$, S1) or positive ($SENT > 0$, S2) sentiment and neutrality below (N1) or above (N2) the median. The rows starting with "Days" report the equally weighted mean excess return of constituent stocks within the time window indicated, with the corresponding t-statistic beneath it. Each panel uses $SENT$ and $NEUT$ from the BERT model shown in the title.

Table 5: Risk-adjusted Returns on Sentiment Portfolios around Stocktwits Post

Panel A: FinBERT

	S1	S2
Days -10 to 0	-0.00740	0.00109
t	-26.87229	4.76398
Days 1 to 7	0.00019	-0.00372
t	0.48901	-13.38112
Days 8 to 14	-0.00050	-0.00323
t	-1.25295	-11.64589
Days 15 to 21	-0.00080	-0.00310
t	-2.04187	-11.26123
Days 22 to 28	-0.00058	-0.00308
t	-1.49065	-11.09403

Panel B: FinBERT-tone

	S1	S2
Days -10 to 0	-0.00582	-0.00084
t	-17.42732	-4.03603
Days 1 to 7	0.00006	-0.00303
t	0.12451	-11.91730
Days 8 to 14	0.00101	-0.00327
t	1.95881	-12.94737
Days 15 to 21	0.00099	-0.00331
t	1.99363	-13.11249
Days 22 to 28	0.00016	-0.00291
t	0.31870	-11.37962

Panel C: Sentiment-xDistil

	S1	S2
Days -10 to 0	-0.00819	0.00072
t	-27.81524	3.26145
Days 1 to 7	0.00055	-0.00355
t	1.28774	-13.42522
Days 8 to 14	0.00014	-0.00328
t	0.33005	-12.30230
Days 15 to 21	0.00017	-0.00334
t	0.40560	-12.55263
Days 22 to 28	-0.00020	-0.00304
t	-0.47244	-11.41460

Stocks are sorted by the sentiment measure ($SENT$) into portfolios every day (Day 0) between November 2, 2020 and August 4, 2023. S1 is the negative sentiment portfolio ($SENT \leq 0$) and S2 the positive sentiment portfolio ($SENT > 0$). The rows starting with "Days" report the equally weighted mean of the risk-adjusted return on constituent stocks in Equation (8) within the time window indicated, with the corresponding t-statistic beneath it. Each panel uses $SENT$ from the BERT model shown in the title.

Table 6: Risk-adjusted Returns on Sentiment-Neutrality Portfolios around Stocktwits Post

Panel A: FinBERT

	S1N1	S1N2	S2N1	S2N2
Days -10 to 0	-0.00815	-0.00519	0.00247	-0.00128
t	-25.75510	-9.28294	8.20049	-3.65032
Days 1 to 7	0.00031	-0.00017	-0.00414	-0.00300
t	0.70728	-0.21014	-11.92445	-6.45080
Days 8 to 14	-0.00036	-0.00091	-0.00391	-0.00204
t	-0.77654	-1.15799	-11.43506	-4.31442
Days 15 to 21	-0.00076	-0.00092	-0.00377	-0.00194
t	-1.69545	-1.14109	-11.16832	-4.08856
Days 22 to 28	-0.00066	-0.00035	-0.00340	-0.00254
t	-1.46634	-0.44844	-9.82777	-5.43122

Panel B: FinBERT-tone

	S1N1	S1N2	S2N1	S2N2
Days -10 to 0	-0.00781	-0.00422	-0.00054	-0.00238
t	-15.29744	-9.57701	-2.36788	-4.80477
Days 1 to 7	-0.00011	0.00020	-0.00313	-0.00251
t	-0.15843	0.30252	-11.34502	-3.84237
Days 8 to 14	0.00097	0.00105	-0.00361	-0.00146
t	1.19833	1.56830	-13.22601	-2.23226
Days 15 to 21	0.00172	0.00040	-0.00356	-0.00201
t	2.25221	0.61159	-13.17304	-2.92647
Days 22 to 28	0.00065	-0.00024	-0.00294	-0.00277
t	0.86367	-0.37151	-10.54013	-4.30400

Panel C: Sentiment-xDistil

	S1N1	S1N2	S2N1	S2N2
Days -10 to 0	-0.00968	-0.00628	0.00188	-0.00084
t	-24.24158	-14.45139	6.15630	-2.69257
Days 1 to 7	0.00150	-0.00067	-0.00406	-0.00287
t	2.76591	-0.99341	-11.50715	-7.18218
Days 8 to 14	0.00133	-0.00135	-0.00389	-0.00247
t	2.20106	-2.15361	-11.07091	-6.02976
Days 15 to 21	0.00054	-0.00029	-0.00389	-0.00261
t	0.95838	-0.44249	-11.35042	-6.21462
Days 22 to 28	0.00025	-0.00077	-0.00355	-0.00236
t	0.43649	-1.17473	-10.12800	-5.76275

Stocks are sorted independently by the sentiment ($SENT$) and neutrality ($NEUT$) measures into portfolios every day (Day 0) between November 2, 2020 and August 4, 2023. Four portfolios are formed as the intersection of negative ($SENT \leq 0$, S1) or positive ($SENT > 0$, S2) sentiment and neutrality below (N1) or above (N2) the median. The rows starting with "Days" report the equally weighted mean of the risk-adjusted return on constituent stocks in Equation (8) within the time window indicated, with the corresponding t-statistic beneath it. Each panel uses $SENT$ and $NEUT$ from the BERT model shown in the title.

Table 7: Panel Regressions on *SENT*

Panel A: FinBERT

	(1)	(2)	(3)	(4)
SENT	0.0468***	0.0530***	0.0588***	0.0655***
	(0.0022)	(0.0030)	(0.0033)	(0.0043)
SENT_m1		-0.0070**	-0.0064*	-0.0106**
		(0.0026)	(0.0027)	(0.0034)
SENT_m5_m2			-0.0035	-0.0100*
			(0.0034)	(0.0046)
SENT_m10_m6			-0.0162***	-0.0194***
			(0.0042)	(0.0053)
SENT_m30_m11			-0.0104	-0.0249**
			(0.0068)	(0.0095)
preSIZE				-0.8373
				(0.9049)
presumDollarVol				8.7819*
				(4.3523)
presumTradeCount				-0.1768**
				(0.0643)
presumShortIntVol				-0.0161
				(0.0140)
preminLastPrice				-0.2915***
				(0.0378)
user_followers_sum				0.2103**
				(0.0791)
user_like_count_sum				1.3952
				(1.9193)
label				0.0444***
				(0.0037)
N	129871	72941	57662	45701
F-stat	701.3701	261.1008	113.1517	79.6156
Adj. R^2				
F-stat (robust)	467.1946	158.1084	68.1554	34.2766
R^2 (Within)	0.0054	0.0072	0.0098	0.0225
AIC				
BIC				

Panel B: FinBERT-tone

	(1)	(2)	(3)	(4)
SENT	0.0260***	0.0280***	0.0296***	0.0289***
	(0.0016)	(0.0020)	(0.0022)	(0.0028)
SENT_m1		-0.0033	-0.0015	-0.0042
		(0.0021)	(0.0022)	(0.0028)
SENT_m5_m2			-0.0018	-0.0019
			(0.0030)	(0.0038)
SENT_m10_m6			-0.0106**	-0.0138**
			(0.0033)	(0.0043)
SENT_m30_m11			-0.0015	-0.0083
			(0.0067)	(0.0084)
preSIZE				-0.9389
				(0.9430)
presumDollarVol				9.3546*
				(4.4478)
presumTradeCount				-0.1864**
				(0.0658)
presumShortIntVol				-0.0176
				(0.0145)
preminLastPrice				-0.2782***
				(0.0394)
user_followers_sum				0.2178**
				(0.0820)
user_like_count_sum				1.3299
				(1.9164)
label				0.0492***
				(0.0038)
N	129871	72941	57662	45701
F-stat	280.2635	100.0900	40.1403	57.1403
Adj. R^2				
F-stat (robust)	277.0931	102.4422	39.1950	26.5120
R^2 (Within)	0.0022	0.0028	0.0035	0.0162
AIC				
BIC				

Panel C: Sentiment-xDistil

	(1)	(2)	(3)	(4)
SENT	0.0128***	0.0254***	0.0336***	0.0500***
	(0.0023)	(0.0025)	(0.0025)	(0.0032)
SENT_m1		-0.0011	-0.0020	-0.0064**
		(0.0018)	(0.0019)	(0.0022)
SENT_m5_m2			-0.0044+	-0.0076*
			(0.0025)	(0.0033)
SENT_m10_m6			-0.0087***	-0.0152***
			(0.0026)	(0.0036)
SENT_m30_m11			-0.0086+	-0.0176**
			(0.0049)	(0.0065)
preSIZE				-0.8353
				(0.9593)
presumDollarVol				8.7910*
				(4.3705)
presumTradeCount				-0.1774**
				(0.0636)
presumShortIntVol				-0.0160
				(0.0139)
preminLastPrice				-0.2756***
				(0.0340)
user_followers_sum				0.2128**
				(0.0804)
user_like_count_sum				1.3697
				(1.9176)
label				0.0421***
				(0.0037)
N	129871	72941	57662	45701
F-stat	148.5123	150.2753	87.6843	87.4212
Adj. R^2				
F-stat (robust)	31.9994	51.1732	39.9924	37.8046
R^2 (Within)	0.0012	0.0042	0.0076	0.0246
AIC				
BIC				

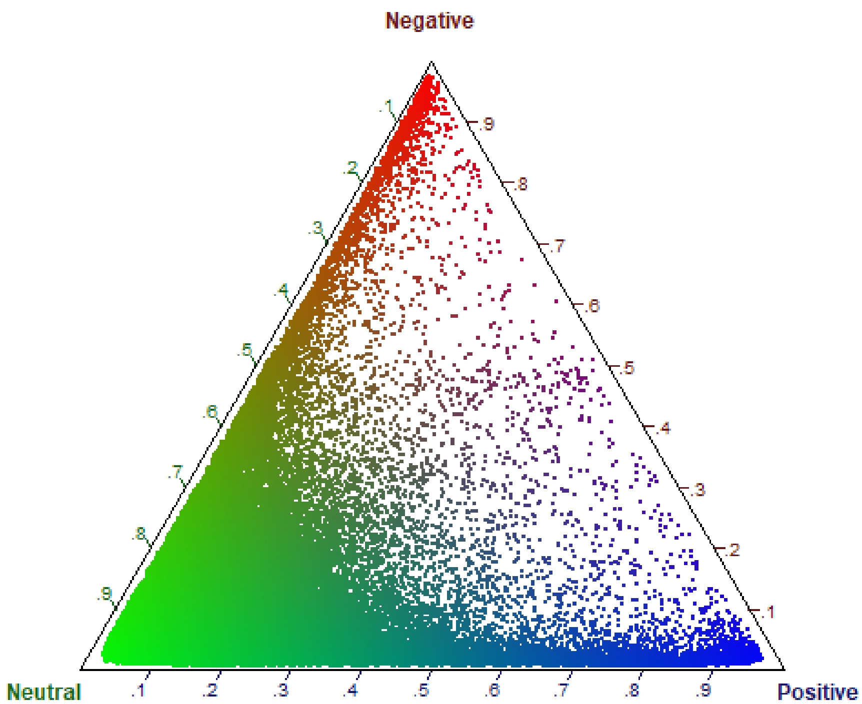

Figure 1: Ternary plot of FinBERT sentiment probabilities.

Figure 2: Distribution of Sentiment Probabilities. This figure plots the prior and posterior densities of the Dirichlet-distributed sentiment probabilities. The Negative and Positive axes represent p_- and p_+, respectively. The neutral probability can be inferred by $p_\circ = 1 - p_- - p_+$ via Equation (1). Panel (A) shows the diffuse prior with concentration parameter $\alpha_0 = (1, 1, 1)$. Panels (B) and (C) depict the posterior densities for ticker symbols VENG on January 23, 2023 and NNMX on January 4, 2022, respectively, where the concentration parameter is computed from Equation (3). These two stocks have the lowest and highest daily excess returns, respectively, in our sample.

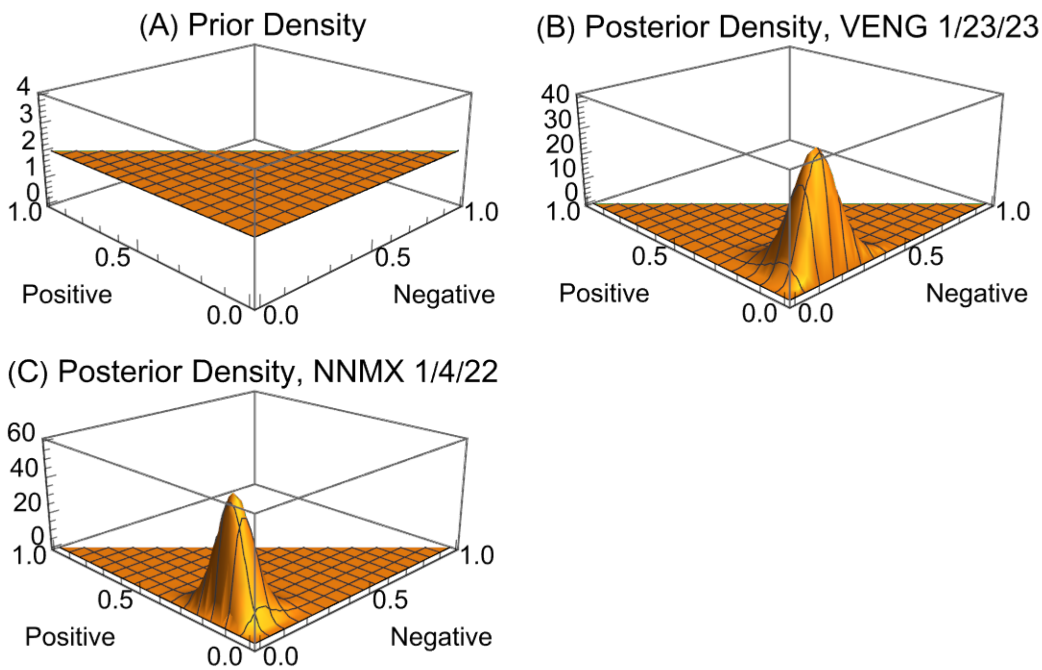

Figure 3: Histograms of Sentiment and Neutrality

Panel A: FinBERT

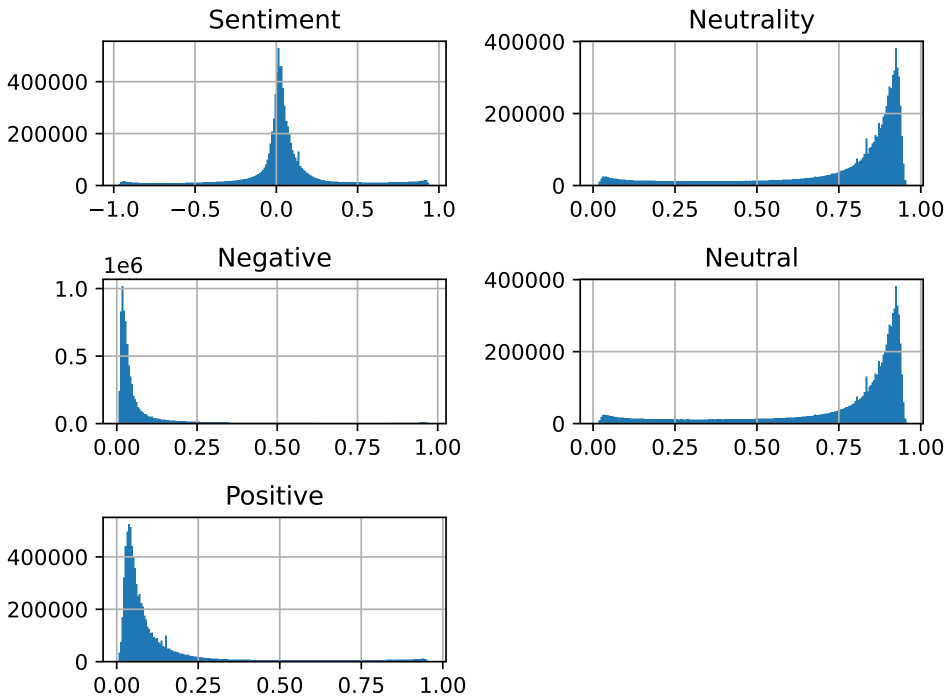

Panel B: FinBERT-tone

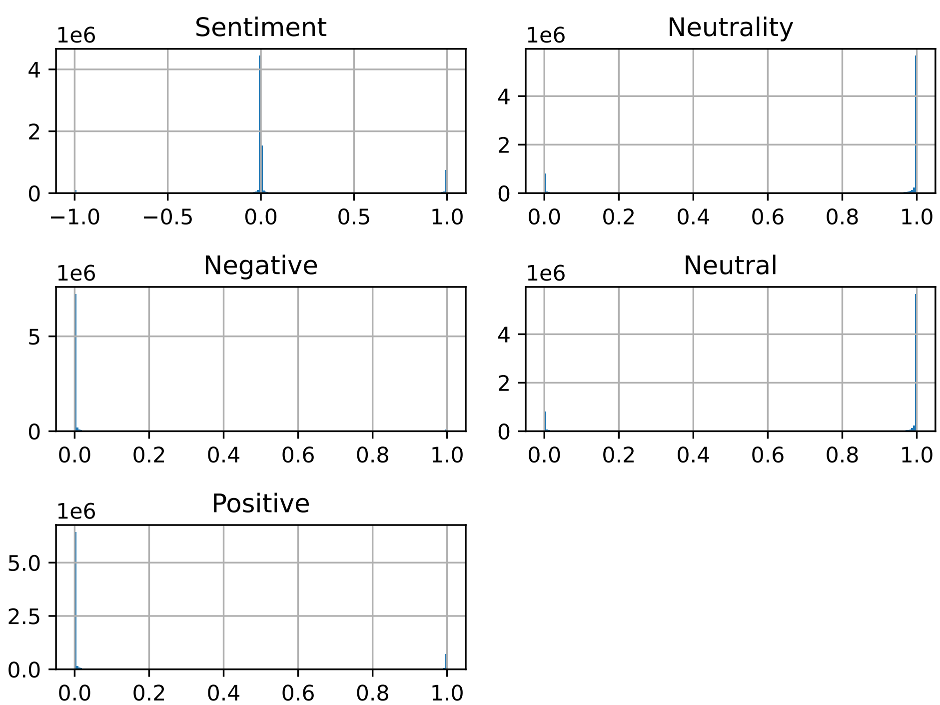

Panel C: Sentiment-xDistil

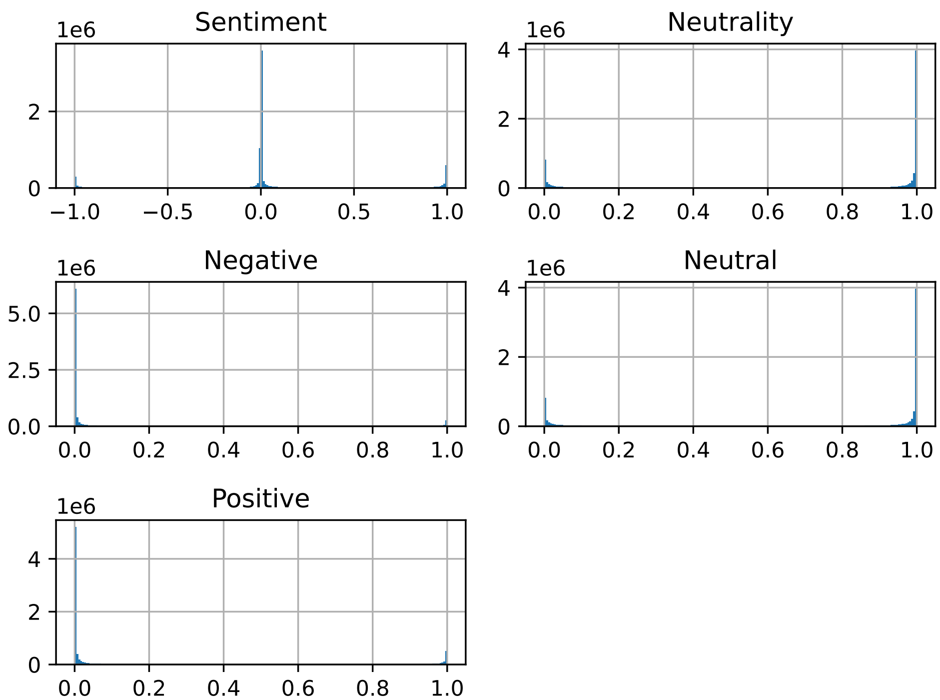

Figure 4: Sentiment Portfolios around Stocktwits Post

Panel A: Sentiment Portfolios

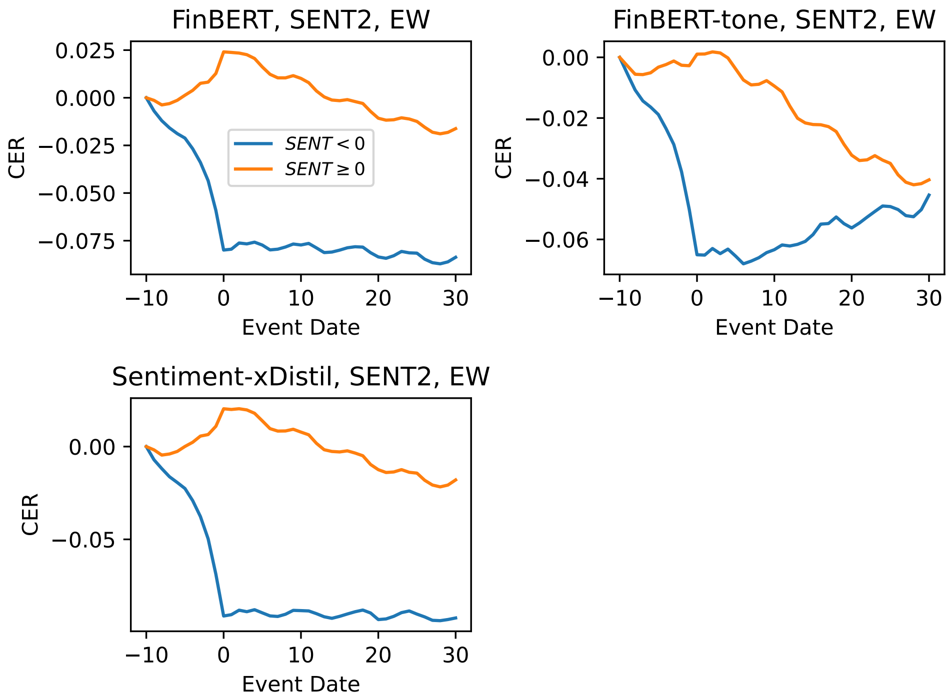

Panel B: Sentiment-Neutrality Portfolios

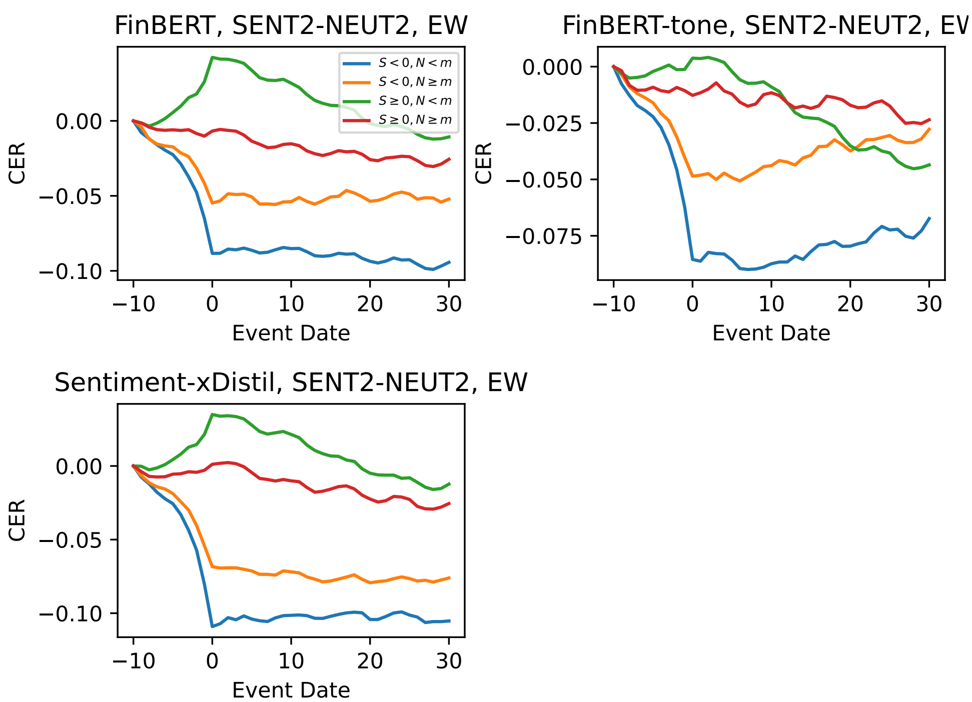

Figure 5: Cumulative Risk-adjusted Return

Panel A: Sentiment Portfolios

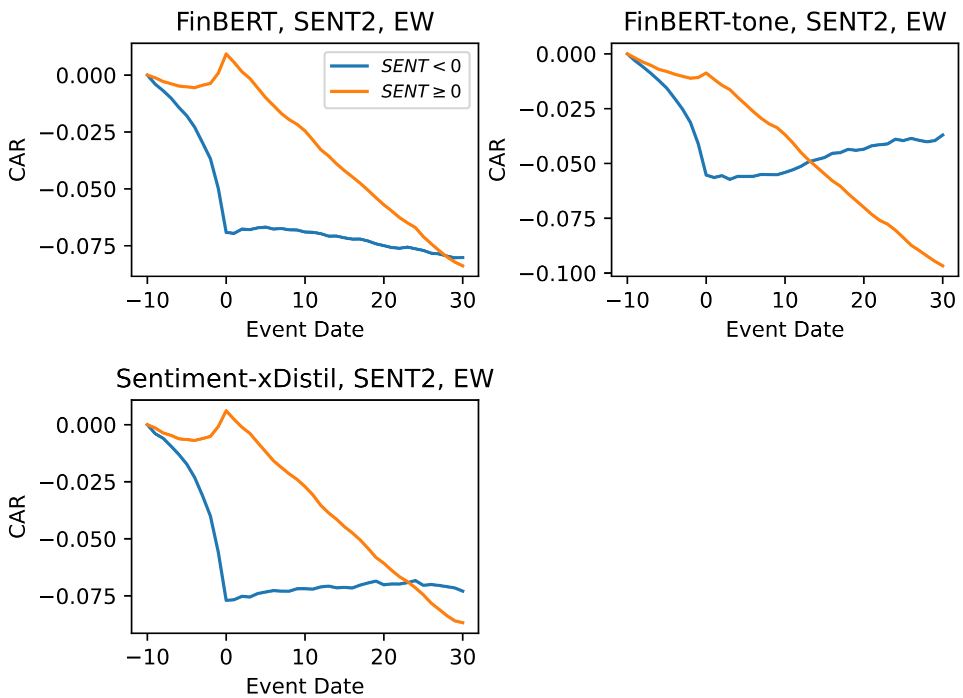

Panel B: Sentiment-Neutrality Portfolios

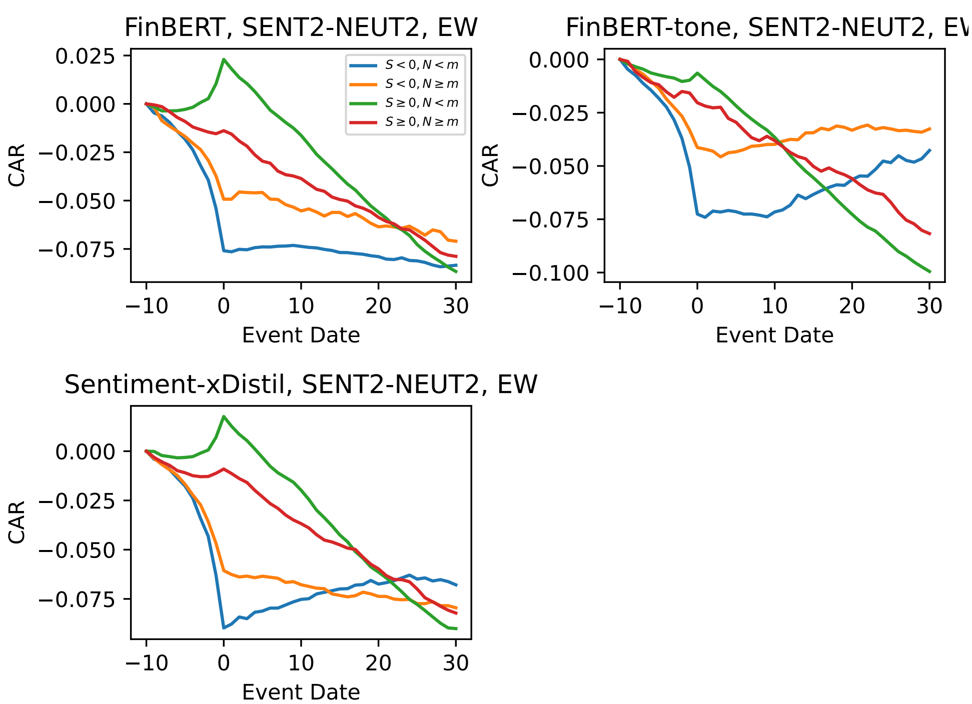

Figure 6: Word Clouds

Figure 7: Word Clouds by Sentiment and Neutrality Levels

Panel A: FinBERT

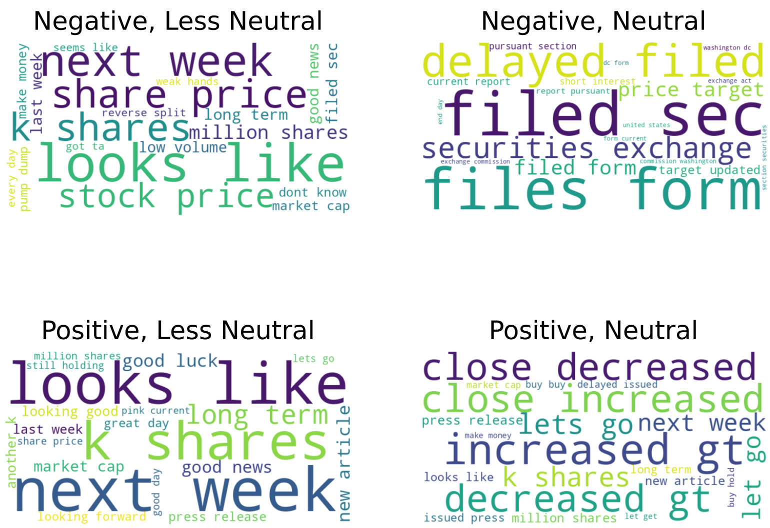

Panel B: FinBERT-tone

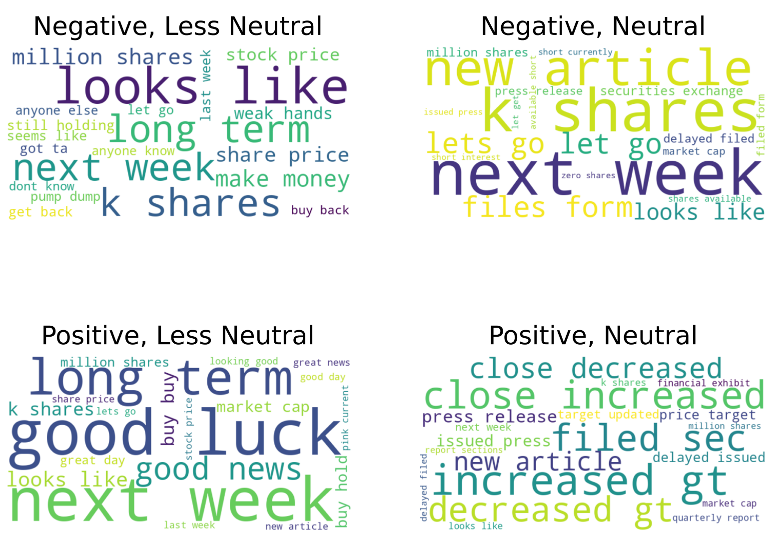

Panel C: Sentiment-xDistil

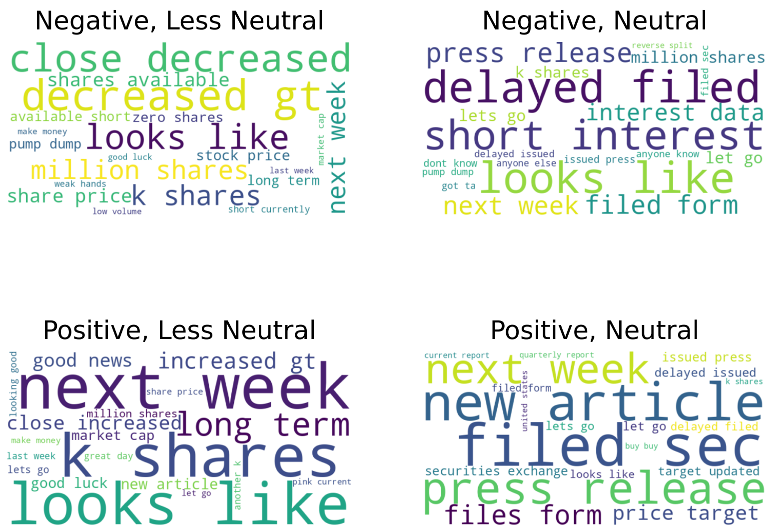

Figure 8: Word Clouds by User Expertise

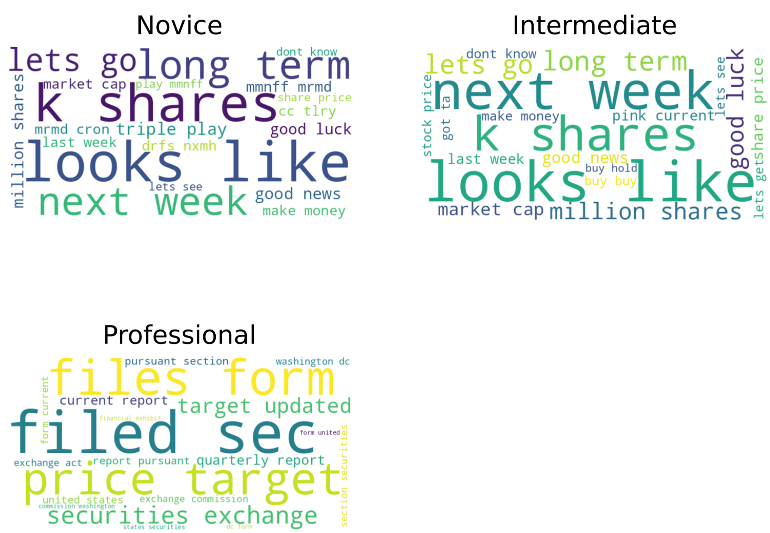

A Appendix

Table A1: Most Frequent Words

1	(('looks', 'like'), 64809)
2	(('next', 'week'), 64060)
3	(('k', 'shares'), 53306)
4	(('new', 'article'), 45680)
5	(('long', 'term'), 44734)
6	(('million', 'shares'), 40070)
7	(('market', 'cap'), 36805)
8	(('press', 'release'), 35883)
9	(('lets', 'go'), 34843)
10	(('let', 'go'), 32000)
11	(('good', 'luck'), 31269)
12	(('good', 'news'), 31248)
13	(('filed', 'sec'), 30240)
14	(('share', 'price'), 28553)
15	(('issued', 'press'), 25721)
16	(('stock', 'price'), 25221)
17	(('increased', 'gt'), 25210)
18	(('close', 'increased'), 25157)
19	(('last', 'week'), 24904)
20	(('decreased', 'gt'), 24697)

This table shows the most frequent words.

Table A2: Most Frequent Words by Sentiment and Neutrality

Panel A: FinBERT

	S1A1	S1A2	S2A1	S2A2
1	((’looks’, ’like’), 18794)	((’filed’, ’sec’), 17745)	((’next’, ’week’), 27791)	((’increased’, ’gt’), 23098)
2	((’next’, ’week’), 12464)	((’files’, ’form’), 15701)	((’looks’, ’like’), 27189)	((’close’, ’increased’), 23055)
3	((’share’, ’price’), 9898)	((’delayed’, ’filed’), 14825)	((’k’, ’shares’), 22372)	((’decreased’, ’gt’), 21974)
4	((’k’, ’shares’), 9351)	((’securities’, ’exchange’), 13515)	((’long’, ’term’), 20762)	((’close’, ’decreased’), 21967)
5	((’stock’, ’price’), 9291)	((’price’, ’target’), 13097)	((’good’, ’luck’), 20352)	((’lets’, ’go’), 19571)
6	((’million’, ’shares’), 9286)	((’filed’, ’form’), 12994)	((’new’, ’article’), 20101)	((’k’, ’shares’), 18369)
7	((’long’, ’term’), 8319)	((’target’, ’updated’), 11320)	((’good’, ’news’), 18616)	((’next’, ’week’), 18284)
8	((’filed’, ’sec’), 7692)	((’current’, ’report’), 7529)	((’market’, ’cap’), 14843)	((’let’, ’go’), 18091)
9	((’good’, ’news’), 7479)	((’pursuant’, ’section’), 7529)	((’another’, ’k’), 11435)	((’press’, ’release’), 16847)
10	((’last’, ’week’), 7314)	((’short’, ’interest’), 7523)	((’looking’, ’good’), 11182)	((’million’, ’shares’), 16131)
11	((’low’, ’volume’), 7147)	((’report’, ’pursuant’), 7522)	((’last’, ’week’), 11029)	((’new’, ’article’), 15273)
12	((’market’, ’cap’), 7054)	((’exchange’, ’commission’), 6810)	((’press’, ’release’), 10997)	((’looks’, ’like’), 14209)
13	((’dont’, ’know’), 6552)	((’united’, ’states’), 6735)	((’great’, ’day’), 10542)	((’issued’, ’press’), 13129)
14	((’make’, ’money’), 6547)	((’exchange’, ’act’), 6697)	((’looking’, ’forward’), 10453)	((’long’, ’term’), 12866)
15	((’pump’, ’dump’), 6299)	((’form’, ’current’), 6642)	((’million’, ’shares’), 10243)	((’delayed’, ’issued’), 12570)
16	((’reverse’, ’split’), 6265)	((’section’, ’securities’), 6584)	((’lets’, ’go’), 9524)	((’buy’, ’buy’), 12431)
17	((’seems’, ’like’), 6112)	((’washington’, ’dc’), 6556)	((’good’, ’day’), 9424)	((’market’, ’cap’), 11440)
18	((’got’, ’ta’), 6008)	((’commission’, ’washington’), 6522)	((’still’, ’holding’), 9155)	((’buy’, ’hold’), 9739)
19	((’every’, ’day’), 5862)	((’dc’, ’form’), 6477)	((’share’, ’price’), 9084)	((’make’, ’money’), 8375)
20	((’weak’, ’hands’), 5549)	((’end’, ’day’), 6365)	((’pink’, ’current’), 8542)	((’let’, ’get’), 8205)

Panel B: FinBERT-tone

	S1A1	S1A2	S2A1	S2A2
1	((’looks’, ’like’), 19617)	((’next’, ’week’), 18004)	((’good’, ’luck’), 28871)	((’increased’, ’gt’), 24067)
2	((’k’, ’shares’), 10523)	((’k’, ’shares’), 16563)	((’next’, ’week’), 27442)	((’close’, ’increased’), 24053)
3	((’long’, ’term’), 10168)	((’new’, ’article’), 16065)	((’long’, ’term’), 26620)	((’filed’, ’sec’), 22508)
4	((’next’, ’week’), 10118)	((’lets’, ’go’), 15504)	((’good’, ’news’), 24863)	((’decreased’, ’gt’), 20948)
5	((’make’, ’money’), 9384)	((’let’, ’go’), 15415)	((’looks’, ’like’), 24315)	((’close’, ’decreased’), 20939)
6	((’million’, ’shares’), 8562)	((’files’, ’form’), 15218)	((’k’, ’shares’), 18947)	((’new’, ’article’), 17298)
7	((’share’, ’price’), 8227)	((’looks’, ’like’), 15033)	((’buy’, ’buy’), 18243)	((’press’, ’release’), 16572)
8	((’stock’, ’price’), 8192)	((’securities’, ’exchange’), 13139)	((’market’, ’cap’), 16402)	((’issued’, ’press’), 14351)
9	((’weak’, ’hands’), 7513)	((’million’, ’shares’), 12663)	((’buy’, ’hold’), 15336)	((’delayed’, ’issued’), 13575)
10	((’still’, ’holding’), 7420)	((’delayed’, ’filed’), 11382)	((’million’, ’shares’), 13725)	((’price’, ’target’), 13231)
11	((’got’, ’ta’), 7170)	((’press’, ’release’), 10273)	((’great’, ’day’), 12105)	((’target’, ’updated’), 11999)
12	((’last’, ’week’), 7156)	((’market’, ’cap’), 8976)	((’looking’, ’good’), 11910)	((’next’, ’week’), 8496)
13	((’seems’, ’like’), 7129)	((’filed’, ’form’), 8815)	((’share’, ’price’), 11731)	((’k’, ’shares’), 7273)
14	((’get’, ’back’), 7128)	((’shares’, ’available’), 8732)	((’good’, ’day’), 11250)	((’quarterly’, ’report’), 6346)
15	((’let’, ’go’), 6842)	((’zero’, ’shares’), 7940)	((’lets’, ’go’), 10952)	((’market’, ’cap’), 6238)
16	((’pump’, ’dump’), 6761)	((’available’, ’short’), 7927)	((’new’, ’article’), 10791)	((’financial’, ’exhibit’), 6120)
17	((’anyone’, ’else’), 6523)	((’let’, ’get’), 7777)	((’last’, ’week’), 10521)	((’looks’, ’like’), 5844)
18	((’anyone’, ’know’), 6432)	((’short’, ’currently’), 7651)	((’great’, ’news’), 10212)	((’million’, ’shares’), 5120)
19	((’buy’, ’back’), 6152)	((’short’, ’interest’), 7531)	((’stock’, ’price’), 9884)	((’report’, ’sections’), 4869)
20	((’dont’, ’know’), 5954)	((’issued’, ’press’), 7418)	((’pink’, ’current’), 9231)	((’delayed’, ’filed’), 4761)

Panel C: Sentiment-xDistil

	S1A1	S1A2	S2A1	S2A2
1	(('decreased', 'gt'), 24693)	(('delayed', 'filed'), 6240)	(('next', 'week'), 33094)	(('filed', 'sec'), 25637)
2	(('close', 'decreased'), 24678)	(('looks', 'like'), 6141)	(('k', 'shares'), 30047)	(('new', 'article'), 20824)
3	(('looks', 'like'), 19631)	(('short', 'interest'), 5525)	(('looks', 'like'), 28989)	(('press', 'release'), 18095)
4	(('k', 'shares'), 11485)	(('press', 'release'), 5151)	(('long', 'term'), 26615)	(('next', 'week'), 16986)
5	(('million', 'shares'), 10511)	(('filed', 'form'), 4781)	(('increased', 'gt'), 25014)	(('files', 'form'), 16344)
6	(('next', 'week'), 9370)	(('next', 'week'), 4610)	(('close', 'increased'), 24961)	(('price', 'target'), 15942)
7	(('share', 'price'), 9322)	(('interest', 'data'), 4374)	(('good', 'news'), 22541)	(('securities', 'exchange'), 14062)
8	(('shares', 'available'), 9257)	(('million', 'shares'), 4179)	(('new', 'article'), 20494)	(('issued', 'press'), 13743)
9	(('stock', 'price'), 8791)	(('lets', 'go'), 4072)	(('good', 'luck'), 19860)	(('delayed', 'issued'), 12768)
10	(('zero', 'shares'), 8520)	(('let', 'go'), 4059)	(('market', 'cap'), 19651)	(('let', 'go'), 12584)
11	(('available', 'short'), 8458)	(('k', 'shares'), 4004)	(('million', 'shares'), 18293)	(('target', 'updated'), 12246)
12	(('long', 'term'), 8255)	(('issued', 'press'), 3775)	(('lets', 'go'), 15952)	(('lets', 'go'), 11587)
13	(('pump', 'dump'), 8113)	(('dont', 'know'), 3738)	(('last', 'week'), 12895)	(('delayed', 'filed'), 10660)
14	(('short', 'currently'), 8006)	(('delayed', 'issued'), 3715)	(('great', 'day'), 12773)	(('looks', 'like'), 10048)
15	(('weak', 'hands'), 7059)	(('got', 'ta'), 3640)	(('another', 'k'), 12273)	(('filed', 'form'), 8820)
16	(('market', 'cap'), 6931)	(('pump', 'dump'), 3546)	(('looking', 'good'), 11938)	(('quarterly', 'report'), 8423)
17	(('make', 'money'), 6554)	(('anyone', 'know'), 3503)	(('share', 'price'), 11895)	(('united', 'states'), 8156)
18	(('good', 'luck'), 6298)	(('filed', 'sec'), 3363)	(('pink', 'current'), 11396)	(('buy', 'buy'), 8030)
19	(('low', 'volume'), 6283)	(('anyone', 'else'), 3117)	(('make', 'money'), 11135)	(('current', 'report'), 7778)
20	(('last', 'week'), 6244)	(('reverse', 'split'), 3040)	(('let', 'go'), 11076)	(('k', 'shares'), 7770)

This table shows the most frequent words by sentiment and neutrality levels as determined by the three BERT methods.

Table A3: Most Frequent Words by User Expertise

	Novice	Intermediate	Professional
1	(('looks', 'like'), 3812)	(('looks', 'like'), 11680)	(('filed', 'sec'), 24327)
2	(('k', 'shares'), 3357)	(('next', 'week'), 10627)	(('files', 'form'), 16155)
3	(('long', 'term'), 2899)	(('k', 'shares'), 8527)	(('price', 'target'), 13869)
4	(('next', 'week'), 2793)	(('long', 'term'), 7746)	(('securities', 'exchange'), 13575)
5	(('lets', 'go'), 2439)	(('lets', 'go'), 7216)	(('target', 'updated'), 12357)
6	(('triple', 'play'), 2022)	(('million', 'shares'), 6512)	(('quarterly', 'report'), 8403)
7	(('million', 'shares'), 1842)	(('good', 'luck'), 6372)	(('current', 'report'), 7667)
8	(('good', 'news'), 1834)	(('market', 'cap'), 5845)	(('pursuant', 'section'), 7610)
9	(('good', 'luck'), 1682)	(('good', 'news'), 4759)	(('report', 'pursuant'), 7537)
10	(('mmnff', 'mrmd'), 1442)	(('share', 'price'), 4468)	(('united', 'states'), 7074)
11	(('mrmd', 'cron'), 1362)	(('pink', 'current'), 4402)	(('exchange', 'commission'), 6833)
12	(('make', 'money'), 1332)	(('last', 'week'), 4401)	(('exchange', 'act'), 6675)
13	(('last', 'week'), 1303)	(('make', 'money'), 3824)	(('form', 'current'), 6654)
14	(('cc', 'tlry'), 1287)	(('dont', 'know'), 3783)	(('washington', 'dc'), 6637)
15	(('market', 'cap'), 1260)	(('got', 'ta'), 3579)	(('section', 'securities'), 6631)
16	(('drfs', 'nxmh'), 1226)	(('stock', 'price'), 3556)	(('commission', 'washington'), 6618)
17	(('share', 'price'), 1224)	(('lets', 'see'), 3526)	(('dc', 'form'), 6573)
18	(('play', 'mmnff'), 1224)	(('lets', 'get'), 3442)	(('financial', 'exhibit'), 6547)
19	(('dont', 'know'), 1193)	(('buy', 'buy'), 3404)	(('form', 'united'), 6330)
20	(('lets', 'see'), 1166)	(('buy', 'hold'), 3234)	(('states', 'securities'), 6320)

This table shows the most frequent words by user expertise.