Wisdom or Whims? Decoding Investor Sentiment, Trading Strategies, and Market Outcomes with AI*

Shuaiyu Chen[†] Lin Peng [‡] Dexin Zhou[§]

Abstract

We leverage rich social media data and large language models (LLMs) to examine the relationship between investor trading strategies, sentiment, and market outcomes. Extracting trading strategies embedded in 96 million social media posts, we find that strategy adoption is heterogeneous and dynamic, with substantial differences in performance outcomes. Our results show that news arrivals decrease users' reliance on technical signals and increase their utilization of fundamental signals. Technical sentiment negatively predicts stock returns, particularly among short-term or inexperienced users, whereas fundamental sentiment positively forecasts returns. Additionally, message sentiment correlates positively with aggregate retail buying, with technical sentiment strongly associated with aggressive buying by Robinhood investors. Our study demonstrates the promise of using AI to understand investor behaviors and their implications for market dynamics.

Keywords: Social finance, Social media, Retail trading, Herding, Technical analysis, Fundamental analysis

*We appreciate suggestions from Zhi Da, Chukwuma Dim, Winston Dou, David Hirshleifer, Byoung-Hyun Hwang, Clifton Green, Marina Niessner, Ron Kaniel, John Shim, Dacheng Xiu, and seminar participants at Baruch College, Purdue University, the University of Notre Dame, Binghamton University, McMaster University, Rochester Institute of Technology, George Washington University, Gen AI Conference, and USC Conference on Social and Behavioral Finance.

[†]chen4144@purdue.edu, Mitch Daniels School of Business, Purdue University.

[‡]lin.peng@baruch.cuny.edu., Zicklin School of Business, Baruch College.

[§]dexin.zhou@baruch.cuny.edu., Zicklin School of Business, Baruch College.

1. Introduction

A growing literature focuses on the role of retail investors in financial markets. In particular, the recent arrival of fintech brokerage platforms and social media sites such as Reddit has stimulated increasing stock market participation by retail investors, raising important questions about the decision-making process of retail investors and their impacts. Past research has provided important insights into retail investors' profitability and their market impacts.[1] Despite these advances, the underlying mechanisms of retail trading—the ways in which they approach stock investing, form their beliefs, and the strategies they employ—remain not well understood.

Characterizing retail investors' beliefs and investment strategies is challenging, as these investors represent a diverse population with heterogeneous socioeconomic backgrounds, information sets, and financial sophistication, among others. Recent literature uses surveys to gain insights into investor beliefs.[2] However, conducting surveys is costly and often does not reflect investors' thought process in real time, which limits the ability of surveys to understand how investors form beliefs and make dynamic trading decisions in response to evolving market conditions.

This paper seeks to bridge the gap in our understanding of retail investors' belief formation and trading by leveraging rich, real-time investor social media data and large language models (LLMs). Our innovative approach enables us to infer the trading strategies that retail investors discuss, identify the factors influencing the adoption of these strategies, and explore their implications for trading patterns and stock returns. We find that investors employ a diverse set of strategies, and the adoption is dynamic, often depending on information environments. These varying

[1]These papers either examine the aggregate order imbalance of retail investors inferred from anonymous TAQ data or use subsamples of retail brokerage account data to shed light on retail investor informativeness. Retail investors have been shown to be either noise traders who are susceptible to behavioral biases (see, e.g., Barber and Odean, 2000; Kumar and Lee, 2006; Barber and Odean, 2008; Barber et al., 2022; Bryzgalova et al., 2023), or those whose orders correctly predict future returns or are consistent with liquidity provision (see, e.g., Kaniel et al., 2008; Kelley and Tetlock, 2013; Boehmer et al., 2021; Welch, 2022).

[2]See, for example, Choi and Robertson (2020), Giglio et al. (2021), Chinco et al. (2022), Liu et al. (2022), Jiang et al. (2024b), and Laudenbach et al. (2024).

adoption decisions result in substantial differences in their informativeness and performance outcomes.

To uncover these dynamics, we apply LLMs to analyze a vast social media dataset consisting of 96 million messages covering approximately 7,800 stocks posted by nearly 840,000 users on StockTwits, a leading investor platform, from January 2010 to June 2023. Recent research shows that this platform captures important information about retail investors' activities (e.g., Cookson and Niessner, 2020; Cookson et al., 2023, 2024c). Using this dataset, our models dynamically classify investor messages into four categories based on the trading strategies they reflect: those referencing technical analysis (TA), those discussing fundamental analysis (FA), messages motivated by other trading strategies (OS), and messages that do not mention any trading strategies (NS).

The analysis reveals that 31% of all StockTwits messages can be identified as using an investment strategy.[3] Among strategy-related messages, investors employ a diverse set of approaches in their analyses. Specifically, 29% of the strategy-related posts reference technical signals, 44% mention fundamental signals, and 28% refer to other strategies, such as mergers and acquisitions (M&A).

Our classification closely matches the self-reported investment approaches and strategies of selected investors, who represent approximately 19% of all users. Investors identifying as technical or momentum traders, as well as those with shorter investment horizons, demonstrate a greater reliance on technical analysis (TA). By contrast, self-identified technical investors tend to use fundamental analysis (FA) less frequently, underscoring distinct preferences between these approaches.

In addition to these cross-sectional patterns, the classification captures significant time-series variation in strategy usage at the individual level. Notably, abnormal firm-level news activity reduces reliance on TA, suggesting that StockTwits users turn to technical analysis when alternative sources of information, such as news, are scarce. Conversely, when firm-level news activity increases, FA usage

[3]The remaining 69% of messages often include memes, catchphrases like "to the moon," or other content that does not explicitly mention specific strategies.

rises, indicating a shift toward fundamental analysis as additional information becomes available.

We then investigate the performances of different investment approaches. We find that bullish sentiment in TA messages is significantly negatively associated with returns on the following day, suggesting that reliance on technical analysis may lead to poor short-term performance. In sharp contrast, bullish fundamental strategy sentiment positively predicts future returns, indicating that fundamental analysis provides valuable information. Messages referencing other strategies also show a negative relationship with subsequent returns, while messages that do not mention trading strategies exhibit no significant relationship with future returns.

We next examine the moderating factors—such as investor sophistication, investment horizons, and external events like the 2021 GameStop (GME) short squeeze—that may influence the relationship between strategy sentiment and future returns. These factors capture important dimensions of investor behavior and provide insights into how market conditions and participant characteristics shape the informativeness of different strategies.

We find that investor sophistication plays a critical role in the overall performance of TA and OS strategies. To measure this, we calculate the fraction of professional investor participation using self-reported professional experience. The results show that during periods of high professional participation, TA and OS strategies become more informative. In contrast, the informativeness of FA messages remains largely unaffected by the level of professional participation. These findings highlight the importance of professional expertise in improving the quality of analysis for certain strategies, particularly those reliant on fundamental signals or other non-TA approaches.

Investment horizons also emerge as an important moderating factor. This analysis is motivated by Cookson et al. (2024a), who demonstrate that investment horizons significantly shape investor sentiment. To investigate this, we identify investment horizons from a subset of messages containing horizon-related language. Our

findings reveal that messages with a long-term focus consistently perform better across all strategies. Conversely, the poor performance of TA and OS strategies is primarily concentrated among messages with a short-term focus. This suggests that the time horizon of an investor fundamentally influences the effectiveness of their chosen strategy, with longer horizons associated with better outcomes.

The 2021 GME short squeeze offers a unique opportunity to analyze the impact of external events on the informativeness of strategy sentiment. This analysis is motivated by Bradley et al. (2024), who show that the GME episode significantly altered the informativeness of investor social media platforms such as WallStreetBets. To assess this effect on StockTwits, we interact strategy-specific sentiments with a post-GME indicator variable. The results indicate that across all strategy types, the informativeness of StockTwits sentiments declined in the post-GME period. Notably, the decline is most pronounced for TA, which appears to have been disproportionately affected. When combined with our earlier findings that professional investors provide more informative TA analysis, these results suggest that the influx of unsophisticated investors during the GME episode led to a decline in the overall informativeness of TA. This pattern underscores how external shocks can amplify the role of investor composition in shaping the quality of strategy-specific sentiment.

Our measures of investor strategies also provide insights into the nuances of the informativeness of retail trades. As discussed earlier, retail order flows have been characterized either as noise trading driven by behavioral biases or as reflecting valuable information not yet incorporated into prices. Our approach allows us to better understand the relationship between investor sentiment, the strategies they employ, and their trades. To this end, we first examine the link between StockTwits sentiment and retail order flows and then decompose the contribution of each strategy to the informativeness of these flows. We also analyze whether message sentiments are associated with herding episodes on Robinhood, which are marked by sharp increases in new buyers for a given stock.

We begin by investigating whether sentiment on StockTwits reflects retail order flow. To address this, we regress retail order imbalance measures (Boehmer et al. 2021; Barber et al. 2023) on sentiments for each strategy category. The results show that sentiment measures across all strategy categories are positively related to both retail order imbalance measures. These findings suggest that StockTwits sentiments capture retail trading activities in real time.[4]

Next, we decompose retail order flow imbalance into components attributable to each trading strategy and a residual component. We find that TA and OS sentiments contribute negatively to the ability of retail order imbalance to predict next-day returns. In contrast, FA sentiment contributes positively to the predictability of future returns. Messages unrelated to trading strategies do not significantly affect the overall informativeness of retail order imbalance. Interestingly, the residual component exhibits significant predictive power for future returns, consistent with retail investors possessing private information about firms that is not directly linked to trading strategies. This finding aligns with the explanations provided by Boehmer et al. (2021) and Barber et al. (2023), and suggest that a certain type of retail investors, for example, those that rely on fundamental signals, may indeed exhibit unique insights about future stock returns.

On the other hand, (Barber et al., 2022) demonstrate that attention-driven retail trades are associated with behavioral biases, as evidenced by the significant price reversals that these trades experience in subsequent periods. To further explore investor strategies behind this phenomenon, we examine the relationship between StockTwits sentiments and retail herding episodes. Following Barber et al. (2022), we define retail herding episodes as periods marked by sharp increases in the number of new Robinhood investors in a given stock. Our analysis reveals that all sentiment measures are positively related to both contemporaneous and future herding episodes. Importantly, TA sentiment exhibits the strongest relationship with these

[4]To reinforce this conclusion, we conduct a placebo test using overnight StockTwits sentiments. We show that StockTwits sentiments from this overnight periods are not significantly related to retail order imbalance accumulated during the normal trading hours. This test underscore the idea that StockTwits sentiments correspond to real-time trading activities of retail investors.

herding measures. These results suggest that a key driver of such behavior is the crowding on technical signals frequently promoted on popular investment-focused social media platforms.

These findings highlight the dual nature of retail trading. On the one hand, retail order flows can reflect private information, such as those contained in fundamental signals, that enhances market efficiency. On the other hand, attention-driven trading, particularly during herding episodes, may lead to inefficiencies and reversals, particularly when driven by TA-based sentiment. By disentangling the contributions of different strategies, our analysis sheds light on the complex dynamics underlying retail trading behavior and its implications for market outcomes.

Our study contributes to several strands of literature. First, it extends the aforementioned literature on retail trading and its informativeness. Recent research has reignited interest in retail investors as zero-cost trading platforms have attracted a large number of new investors to financial markets (e.g., Barber et al., 2022; Welch, 2022; Eaton et al., 2022). However, the investment approaches and strategies employed by these retail investors have not been well understood. Our research advances this literature by directly extracting investor strategies from their own words and linking these strategies to their trades. We provide the novel finding that retail order flow informativeness is contingent on the types of dominant strategies as reflected on popular social media platforms and that retail investors' herding on technical signals creates profitable opportunities for sophisticated traders.

We also contribute to the literature on the understanding of investors' utilization of different trading strategies. Early literature indicates the effectiveness of fundamental analysis and technical analysis.[5] However, it is unclear if retail investors are able to leverage the investment strategies documented in these studies.

[5]For earlier studies on fundamental analysis, see, for example, Porta et al. (1997) and Abarbanell and Bushee (1997). Early work on technical analysis includes Brown and Jennings (1989); Jegadeesh (1991); Brock et al. (1992); Jegadeesh and Titman (1993a); Blume et al. (1994); Lo et al. (2000); George and Hwang (2004). Building on previous studies that use price patterns to predict future returns, an emerging set of new studies reexamines the potential effectiveness of using price and volume patterns to predict future returns (e.g., Han et al., 2013, 2016; Jiang et al., 2023; Murray et al., 2024).

We provide new evidence on the heterogeneity of investors' ability to use these investment approaches.

Another strand of recent research underscores the role of social networks in shaping retail investors' decisions and uses social media data as a lens to infer investor decision-making.[6] A closely related paper is that of Cookson et al. (2024b), who finds that user sentiment on investor social media platforms positively predicts one-day-ahead returns. Our key contribution is identifying significant heterogeneity in the informativeness of social media sentiment—while sentiment from posts focused on fundamental analysis *positively* predicts future returns, sentiment from technical analysis posts *negatively* predicts them, with this predictability lasting up to a week. We also show that these sentiments are strongly associated with investor herding episodes and the profitability of sophisticated investors, suggesting that social media sentiments play a potentially important role in shaping equilibrium prices and trading.

Finally, our study adds to an increasing number of papers that use LLMs to answer economics and finance questions (e.g., Korinek, 2023).[7] We show that such tools, when applied to rich social media data, provide powerful inferences that help us better understand investors' decision-making. Our paper also illustrates a novel, relatively fast, and cost-effective way of implementing LLMs—instead of purely relying on cutting-edge LLMs, one can first generate useful examples using state-of-the-art (SOTA) LLMs and then use these examples to fine-tune a smaller language

[6]Several studies utilizing data from online social networks have shown that the posting activity and message quality of retail investors help predict stock returns and trading volume (e.g., Antweiler and Frank, 2004; Chen et al., 2014). These studies also highlight how investor disagreement and echo chambers influence belief formation (e.g., Giannini et al., 2018, 2019; Cookson and Niessner, 2020; Cookson et al., 2023), how the dissemination of informative content can be affected (e.g., Chen and Hwang, 2022; Farrell et al., 2022; Bradley et al., 2024), and the role of investor horizon differences (Cookson et al. 2024a). Furthermore, there is growing interest in the skill and role of influencers in these networks (e.g., Coval et al., 2021; Kakhbod et al., 2023; Hirshleifer et al., Forthcoming). The review of social media and finance of Cookson et al. (2024c) summarizes this emerging line of research.

[7]For example, Jiang et al. (2024a) and Lopez-Lira and Tang (2023) use LLMs to predict future returns. Li et al. (2023) extract corporate culture from analyst reports. Jha et al. (2024) extract information related to corporate investments, and Eisfeldt et al. (2023) investigate which jobs are more replaceable with the advent of GPTs. Huang et al. (2024) use LLMs to examine the narratives on investor social media.

model.

2. Data

Our sample includes common stocks (CRSP share codes 10, 11, and 12) traded on the NYSE, AMEX, and NASDAQ from January 2010 through June 2023. We obtain investor social media data from StockTwits, stock market data from CRSP, accounting data from Compustat, retail market order data from TAQ, Robinhood user account data from RobinTrack, and financial news data from RavenPack.

2.1. StockTwits Data

StockTwits is a leading social media platform dedicated to retail investors, allowing users to share opinions and exchange ideas about stocks, ETFs, and cryptocurrencies. Similar to Twitter, StockTwits users post short messages, initially limited to 140 characters until May 8, 2019, when the limit expanded to 1,000 characters. A distinguishing feature of StockTwits is its focus on financial markets, with users employing "cashtags" (e.g., $TSLA) to indicate specific ticker symbols mentioned in their posts.

We collect comprehensive message-level data using the StockTwits API, covering 169,509,106 posts from 978,071 users related to 15,232 tickers (including stocks, ETFs, and closed-end funds) between January 2010 and June 2023.[8] At the message level, our data include timestamps, textual content, and user-provided sentiment labels ("bullish" or "bearish") when available. Additionally, we obtain user-level biographical characteristics self-reported by StockTwits users, including investment style (Technical, Momentum, Fundamental, Value, Growth, or Global Macro), investment horizon (Day Trader, Swing Trader, Position Trader, or Long-Term Investor), and trading experience (Novice, Intermediate, or Professional).

[8]The StockTwits API documentation is available at `https://firestream-portal.stocktwits.com/documentation/stream`.

Following Cookson and Niessner (2020) and Cookson et al. (2024b), we apply several filters to ensure message validity and to focus explicitly on content generated by human users discussing publicly traded companies. Specifically, we first retain only messages explicitly referencing exactly one ticker symbol. Second, we exclude users posting more than 1,000 messages on a single day and remove messages sourced from third-party platforms, as these typically redistribute financial news or involve algorithmically generated content. Finally, we require both the user identifier and username fields to be non-missing. After merging these filtered messages with our CRSP stock universe, the resulting final dataset includes 96,095,345 messages from 840,846 unique users covering 7,834 stocks from January 2010 through June 2023.

2.2. Extracting Information from StockTwits Texts

2.2.1. Identifying Trading Strategies From Messages

Leveraging large language models, we first decipher the strategies conveyed in the message content. Specifically, we identify messages related to technical, and fundamental strategies, two important types of investment strategies. We also attempt to identify messages that use other investment strategies. We use the same procedure to accomplish these three classification tasks. First, we leverage the cutting-edge large language model from OpenAI, GPT-4, to identify if a message contains an investment strategy (or TA/FA).[9] As with many social media messages, those on StockTwits tend to be short, with many abbreviated and colloquial words, and with many non-standard spelling. Thus, it is difficult to identify trading strategies purely based on a dictionary. Moreover, given that trading strategies are highly diverse, identifying messages containing trading strategies can be a highly challenging task. Additionally, GPT-4 also provides responses that are highly similar to humans.[10]

[9]Specifically, we use the GPT 4-Turbo model (gpt-4-0125-preview endpoint).

[10]**?** show that GPT-4 makes judgments that are consistent with humans. Moreover, the overlap between GPT-4 and human responses is at the level of humans.

While we find that GPT-4 has an excellent ability to identify these strategy-related messages, and those identifications tend to align with our own judgment, it is infeasible to use GPT-4 to classify all the messages in our sample due to the limited throughput and high costs. Thus, we use the examples generated by the cutting-edge GPT-4 model to fine-tune a smaller classification model.[11]

Next, we illustrate our classification procedure by identifying TA-related messages. We first randomly sample $20,000$ messages from the sample of messages.[12] We then ask the GPT-4 to determine whether the message entails technical trading using the following prompt:

> You have a deep understanding of the language of social media and financial markets. Please analyze the message from an investor social media platform. Please parse the message along two dimensions. 1) Presence of technical analysis (0=no, 1=possibly, 2=likely). 2) if technical analysis is used, what is the technical indicator? (output the indicator or "" if you cannot locate it. If multiple signals exist, please separate by a comma) Output in JSON format: {"technical_analysis":, "technical_indicator": }.

We collect GPT's response for the $20,000$ messages. Table A.1 in the appendix provides a sample of positive and negative responses by GPT. Then, we use these responses to fine-tune a BERT model (henceforth TechBERT) to provide a prediction of whether a message uses technical analysis. Through cross-validation, we find that the fine-tuned TechBERT model can achieve an F1 score of 0.83, which indicates a high level of performance. Since BERT has a drastically smaller parameter count, we are able to run this model locally to provide a probabilistic prediction of whether each message contains technical trading.[13]

[11] This approach, first proposed by Hinton (2015), is widely known as knowledge distillation (KD) in the machine learning literature. See Gu et al. (2023) for a more recent review of this approach and its applications in LLMs.

[12] To achieve a more balanced sample, $10,000$ messages are sampled from users with a self-declared technical investment style and the other $10,000$ are from the other groups.

[13] We use the bert-base-uncased model, which has 110 million parameters. While it has a large

We visualize TechBERT's classification results (Technical Analysis Intensity) in Figure 1 and Figure 2 across investor types depending on their self-labeled investment approaches and horizons, respectively. We find that TechBERT's classification prediction exhibits some desirable properties, as most of the messages fall either in the low probability region (i.e., $< 5\%$) or high probability region (i.e., $> 95\%$). This result shows that TechBERT's prediction is quite unambiguous.

Figure 6 presents word cloud plots to exhibit the high-frequency unigram and bigram in the TA messages, respectively. In the unigram plot (Panel A), there are several striking patterns. First, technical messages often contain analyses of charts, consistent with the finding in Jiang et al. (2023). Second, we can see many familiar technical terms, such as resistance, support, and gap. In the bigram plot (Panel B), besides other common terms referring to technical signals, we also find terms related to horizons, such as short-term and next week.

We follow the same approach to identify the messages that contain an investment strategy and messages containing fundamental analysis. We only need to revise the prompt that we feed into GPT-4. Then, we fine-tune specialized BERT models to help identify fundamental analysis-related messages. For fundamental analysis, we use the following prompt:[14]

> You have a deep understanding of the language of social media and financial markets. Please analyze the message from an investor social media platform. Please parse the message along two dimensions. 1) Presence of fundamental analysis (0=no, 1=possibly, 2=likely). 2) If fundamental analysis is used, select one of the following 15 topics that is most relevant: "acquisitions-mergers","analyst-ratings","assets","bankruptcy", "credit","credit-ratings","dividends","earnings","equity-actions","investor-

parameter count, it is relatively smaller compared to GPT-4, which is rumored to have 1.7 trillion parameters. BERT has established itself as a nimble yet highly capable tool for many natural language processing tasks, including classification (Devlin et al., 2018). González-Carvajal and Garrido-Merchán (2020) show that BERT achieves superior performance compared to traditional natural language processing tools that do not rely on deep learning.

[14]Our fundamental specific topics are formed based on the topic classifications of financial news articles provided by Ravenpack.

relations","labor-issues","marketing","price-targets","products-services", "revenues". Output in JSON format: {"fundamental_analysis":, "fundamental_topic":}.

Similar to TA strategies, we visualize the BERT classification results from fundamental messages in Figure 4 and Figure 5 for self-labeled investment approaches and horizons. We also present word cloud plots in Figure **??** to exhibit the high-frequency unigram and bigram in the FA messages, respectively.

For overall strategy, we use the following prompt:

> You have a deep understanding of the language of social media and financial markets. Please analyze the message from an investor social media platform. Please parse the message along two dimensions. 1) Presence of investment strategy (e.g., technical analysis, fundamental analysis, event-driven strategy, arbitrage strategy). If true, please answer 1, otherwise 0. 2) if a strategy is identified, please specify the strategy Output in JSON format: {"has_strategy":, "strategy_type": }.

We then classify a message as Other Strategy (OS) if it is strategy related but not about either technical or fundamental analysis.

Overall, we find that approximately one-third of all messages involve discussions related to trading strategies. TA, FA, OS related messages comprise roughly 28%, 44%, 28% of the strategy-related messages, respectively.

2.2.2. Sentiment Classification

We follow the spirit of Cookson and Niessner (2020) to assign a sentiment score to StockTwits messages with missing self-reported sentiments.[15] Specifically, we randomly select $100,000$ messages with a self-declared bullish or bearish label. Then, we finetune the BERT model based on the randomly selected messages.[16] The clas-

[15] In StockTwits, users have the option to declare sentiment when posting a message. However, not all messages contain the self-declared sentiment flag.

[16] We also adopt the maximum entropy approach used in Cookson and Niessner (2020) to impute missing sentiment at the message level. Our main findings hold quantitatively and qualitatively based on this alternative sentiment imputation method.

sifier delivers a probabilistic prediction on whether a given message is bullish, and we apply this classifier to messages without sentiment labeling as bullish or bearish.[17]

We aggregate sentiments for each firm and investment strategy category. Following Cookson et al. (2024b), sentiment is defined as the normalized difference between bullish and bearish messages:

$$Sentiment_{i,t} = \frac{N_{i,t}^{Bullish} - N_{i,t}^{Bearish}}{N_{i,t}^{Bullish} + N_{i,t}^{Bearish}}.$$

These sentiments are reported in Table 1 Panel B. We find that sentiments of all four strategies (TA, FA, OS, and NS) have similar standard deviations.

2.3. Other Variables

The list of other stock variables and firm characteristics, as well as their construction, are listed in Table A.2. Panel B in Table 1 reports summary statistics in the firm-day sample, including StockTwits sentiment based on each category, attention, order imbalance measures, and other firm-level characteristics.

3. Determinants of Retail Trading Strategy Usage

In this section, we examine the factors influencing StockTwits users' adoption of different investment strategies. These analyses serve two primary objectives. First, they validate our LLM-based approach for classifying trading strategies at the message level. Second, they provide detailed insights into how trading strategy usage varies across investor characteristics, stocks, and over time.

We begin by investigating whether the strategies inferred from individual messages align with users' self-reported profiles, such as investment approach, holding horizon, and experience. Our findings indicate that the reliance on technical

[17] We use the messages with a bullish/bearish flag that are not in the training sample to conduct model validation. Our classifier achieves an F1 score of 0.9, which indicates the high accuracy of our model.

analysis (TA) is substantially higher among investors whose self-reported characteristics are typically associated with technical trading. For example, when categorizing messages by user-declared investment approaches, Figure 1 reveals that users identifying as technical or momentum traders exhibit greater usage of technical analysis relative to those identifying as fundamental, value, growth, or global macro investors. Consistent with this observation, Figure 4 shows that investors who self-report using fundamental, value, and growth rely more heavily on fundamental analysis (FA) in their StockTwits messages. Furthermore, as illustrated in Figures 2 and 5, our BERT models classify a higher proportion of TA-related (and a lower proportion of FA-related) messages among self-declared day traders and swing traders compared to position and long-term investors.

However, these figures also highlight substantial heterogeneity in investment strategies among users within the same self-reported profiles, which typically remain unchanged since users registered their Stocktwits accounts. To formally analyze these variations in the usage of various strategies, we estimate a series of panel regressions at the message level, in which we restrict to Stocktwits users with available biographical information.[18]

Our regression is specified as follows:

$$Usage_{i,j,t,n}^{type} = \beta_1 X_j^{investor} + \beta_2 Y_{i,t}^{stock} + \beta_3 Z_{i,j,t,n}^{message} + FEs + \epsilon_{i,j,t,n},$$

where $Usage_{i,j,t,n}^{type}$ is an indicator variable set to one if message n, posted by investor j about stock i on day t, is classified by our BERT models into one of the trading strategy categories: TA, FA, or Other Strategy (OS). The vector $X_j^{investor}$ includes indicators representing investors' self-reported biographical characteristics. Specifically, it captures self-reported investment approaches (e.g., *Technical Investor* equals one for the investor identifies with a technical or momentum style), investment horizons (e.g., *Long-Term Investor* equals one if the investor declares a long-term holding

[18]This requirement reduces the size of our message-level sample (see Panel 1 of Table **??** in the appendix).

horizon), and experience levels (e.g., *Novice* and *Professional*). The vector $Y_{i,t}^{stock}$ incorporates firm-day characteristics, such as abnormal turnover and abnormal volume of news articles (sourced from RavenPack), following Barber et al. (2022). These measures help us capture time variations in information flows and investor attention toward specific stocks. Message-specific attributes are represented by the vector $Z_{i,j,t,n}^{message}$, which includes the length of each message and the frequency of technical and fundamental terminology derived via a bag-of-words approach using the lists of technical and fundamental related words from Table II in Cookson and Niessner (2020). We apply a TF-IDF weighting scheme to emphasize terms specific to certain messages and down-weight those appearing commonly across many messages. Finally, we include an extensive set of fixed effects (FEs), such as investor, stock, date, and their interactions, to account for unobserved heterogeneity. These fixed effects also allow us to quantify the variations in retail investors' trading strategy usage that are attributable to factors beyond our observed variables.

The regression results are reported in Table 2. Panel A presents the determinants of technical analysis usage ($Usage^{TA}$). In column (1), our primary explanatory variables are indicators of investors' self-reported investment approaches. The results indicate that self-reported investment styles significantly influence message-level strategy usage. For instance, investors identifying as technical traders ("Technical" or "Momentum") exhibit a 10.3 percentage point higher likelihood of using TA in their messages relative to investors reporting fundamental-related approaches ("Fundamental," "Value," and "Growth"). However, the relatively low R-squared (approximately 2.2%) suggests substantial heterogeneity, potentially arising from inaccuracies in self-reported investor profiles or considerable time variation in strategy adoption.

In column (2), we introduce additional investor characteristics, including investment horizons and trading experience. The estimates reveal several notable patterns. Investors with a short-term focus ("Day Trader" or "Swing Trader") are significantly more likely to adopt technical analysis compared to the benchmark

group ("Position Trader"), whereas long-term investors exhibit a lower propensity for TA usage. In addition, self-reported professional investors are 3.1 percentage points more likely to utilize TA compared to the benchmark group ("Intermediate"), whereas novice investors are significantly less likely to do so. These findings suggest familiarity and experience with technical indicators may drive observed systematic differences, as seasoned investors may be more explicit and detailed when discussing technical reasoning in their posts.

Column (3) further controls for stock-level and message-specific characteristics. We find that abnormal turnover and abnormal news flow negatively correlate with TA usage, indicating that investors rely more heavily on technical analysis when other firm-specific informational signals (e.g., earnings, product services, and analyst recommendations) are less prevalent. Additionally, we incorporate bag-of-words (BoW) TF-IDF scores for technical and fundamental terminology, using word lists from Cookson and Niessner (2020). We observe that messages scoring higher in technical vocabulary exhibit significantly increased TA usage, while fundamental vocabulary scores are negatively associated with TA usage.

In column (4), we add stock and date fixed effects, controlling for variations in TA adoption associated with different stocks and time periods. Intuitively, investors might rely more heavily on technical analysis when analyzing small-cap growth stocks with limited information, such as emerging technology firms lacking positive earnings or dividends. Including date fixed effects is motivated by Cookson et al. (2024b), who document that the StockTwits platform's informativeness changed following its 2019 expansion of message-length limits from 140 to 1,000 characters. The inclusion of these fixed effects moderately increases the R-squared from 6.3% in column (3) to 8.4%, suggesting that unobserved stock- and time-specific factors are not the primary determinants driving the adoption of TA.

In addition to stock and date fixed effects, column (5) introduces investor fixed effects to capture investor-level preferences that are not fully captured by self-reported investor profiles. Including investor fixed effects substantially increases

the R-squared from 8.4% to 21.3%, highlighting investor heterogeneity as the single most critical factor in determining the usage of TA. This finding underscores that relying solely on self-reported biographical information available on social media platforms is insufficient to fully understand retail investors' strategy choices.

Finally, in column (6), we impose stock×investor fixed effects to account for the possibility that investors tailor their strategy choices differently across individual stocks. This more stringent specification further increases the R-squared to 28.7%, indicating considerable variation in TA-related strategy adoption even within the same investor-stock pairs. Collectively, these results emphasize the substantial value provided by message-level analysis in understanding the reliance of retail investors on technical analysis.

In Panel B, we repeat our analyses using FA usage ($Usage^{FA}$) as the dependent variable. We find that investors who self-report technical or momentum styles are less likely to adopt fundamental analysis, whereas investors identifying with fundamental-related approaches exhibit significantly higher FA usage. Long-term investors also demonstrate a greater propensity to post FA-related messages. Additionally, professional investors are more inclined to incorporate fundamental analysis in their posts, suggesting a positive association between experience and the likelihood of employing fundamental language.

In contrast to TA, however, FA usage positively correlates with abnormal news flow, indicating that investors rely more on fundamental analysis when firm-specific news and information are abundant. We also find longer messages have a higher probability of containing FA content.

As with TA usage, the inclusion of user fixed effects substantially increases the regression R-squared by about 10%, suggesting that investor-specific preferences largely drive the variation in investors' FA usage, beyond their self-reported characteristics with investment approach, horizon, and experience.

In Panel C, we examine the usage of other strategies ($Usage^{OS}$), which represent investment approaches unrelated to TA or FA as identified by our BERT models.

The patterns for OS exhibit similarities to those of TA: OS usage is more prevalent among self-declared technical, short-term, and professional investors. Furthermore, OS-related messages are more likely to include technical terms and less likely to feature fundamental words. However, unlike TA, OS usage positively relates to abnormal turnover, suggesting that these strategies may be associated with heightened investor attention and increased market liquidity.

4. Performances of Retail Investment Strategies

While extensive literature has investigated retail investor behavior (e.g., Barber and Odean, 2013; Boehmer et al., 2021; Barber et al., 2023), relatively few studies have explored the heterogeneity across investors' trading strategies and their associated performance outcomes. Existing research provides evidence regarding overall retail investor performance by analyzing either aggregate market order flows or individual trading accounts. However, we still have limited insights into how retail investors adopt specific investment strategies and how these strategies ultimately affect their investment returns. In this section, we address this gap by evaluating the performance of distinct retail investment strategies identified through LLM-based methods applied to StockTwits messages.

4.1. Daily Return Predictability

We first examine whether StockTwits sentiment associated with each type of strategy predicts stock returns over the subsequent trading day. Specifically, we follow Cookson et al. (2024b) to estimate the following panel regression at the stock-day level:

$$Return_{i,t+1} = \beta_1 Sentiment_{i,t}^{type} + \beta_2 Attention_{i,t} + \gamma \boldsymbol{X_{i,t}} + \delta_t + \epsilon_{i,t+1},$$

where the dependent variable, $Return_{i,t+1}$, represents the stock return for stock i

on the next trading day ($t+1$). The key explanatory variables, $Sentiment_{i,t}^{type}$ (with $type \in \{\text{TA, FA, OS, NS}\}$), capture the sentiment embedded within messages categorized into four strategy types: technical analysis (TA), fundamental analysis (FA), other strategies (OS), and non-strategy-related messages (NS). Investor attention to stock i on day t ($Attention_{i,t}$) is defined as the percentage share of StockTwits messages about the stock relative to the total number of messages posted across all stocks on the same day: $\frac{\#Messages_{i,t}}{\sum_i \#Messages_{i,t}} \times 100$. The vector of control variables ($X_{i,t}$) includes log market capitalization, log book-to-market ratio, asset growth, gross profitability, analyst coverage, log institutional ownership, the maximum daily return (MAX) in the previous month, abnormal turnover, abnormal news volume, and lagged returns over the previous five trading days. All regressions include trading-day fixed effects, and standard errors are clustered by trading day.

The regression results are presented in Table 3. In columns (1) through (4), we separately examine sentiment from each investment-strategy category. Our regression results reveal significant heterogeneity in predictive ability across these categories. First, TA-related sentiment is significantly and *negatively* associated with next-day returns, indicating that stocks receiving predominantly bullish technical messages subsequently underperform. Specifically, stocks with the most bearish TA sentiment ($Sentiment^{TA} = -1$) outperform those with the most bullish TA sentiment ($Sentiment^{TA} = 1$) by approximately 3.2 bps (1.6×2) on the following trading day.

In sharp contrast, FA-related sentiment significantly and *positively* predicts the next-day returns: stocks with the most bullish FA sentiment outperform those with the most bearish sentiment by 2.6 bps (1.3×2).

Similar to TA sentiment, sentiment associated with other strategies (OS) negatively predicts returns for the subsequent trading day. In contrast, sentiment from non-strategy-related (NS) messages does not significantly forecast future returns. Column (5) presents the results of a horse-race regression, simultaneously including all four sentiment variables. The findings remain consistent, reflecting low

correlations among sentiment measures across different strategy types and highlighting the robustness of our empirical results.

4.2. Return Predictability at Longer Horizons

We next extend our analysis to examine the predictive power of sentiment measures at longer horizons, motivated by evidence suggesting that investors employing different strategies may differ substantially in their investment horizons (e.g., Cookson et al., 2024a). To evaluate this possibility, we re-estimate regression (4.1), using cumulative stock returns over three extended prediction windows: days $t+1$ to $t+5$, $t+6$ to $t+10$, and $t+11$ to $t+15$. The regression results for these horizons are reported in Table 4.

Panel A shows that the return predictability documented in Table 3 for all three strategy-related sentiments (TA, FA, and OS) is strongest in the shortest window (t+1 to t+5). However, we also find evidence that TA sentiment significantly predicts returns in the intermediate window ($t+6$ to $t+10$), while the predictive power of OS sentiment extends even further, into the longest window ($t+11$ to $t+15$). Interestingly, FA sentiment positively predicts returns only in the shortest horizon, with no evidence of subsequent reversal. This pattern suggests that retail investors employing fundamental strategies may possess timely information regarding near-future news releases or announcements. Taken together, these results demonstrate substantial differences in informativeness across investment strategies discussed on StockTwits: sentiment from FA messages consistently predicts positive future returns, whereas sentiment from TA and OS messages tends to negatively forecast returns, potentially generating losses for investors following these strategies.

Given recent evidence from Cookson et al. (2024a) on heterogeneity in predictive horizons revealed from investor pitches posted on the Motley Fool, we further examine the informativeness of sentiments of various strategies based on implied predictive horizons within messages. Unlike the Motley Fool data studied by Cookson et al. (2024a), StockTwits messages do not explicitly indicate predictive horizons.

To address this empirical challenge, we train a BERT model to classify messages according to their implied predictive horizons into four categories: daily, weekly, long-term, and messages lacking any identifiable horizon. Approximately 70% of messages lack an explicit horizon, leaving roughly 30% of the original message sample available for this analysis.[19] In the set of horizon-classified messages, a large fraction of messages emphasize the daily horizon, consistent with the short-term orientation observed in our main analyses. However, a meaningful proportion of messages also targets longer horizons. To assess how predictive horizons affect strategy performance, we partition messages by both strategy type and horizon category, constructing sentiment measures for each subgroup, and then examining their relationships with future returns.

The results, reported in Panel B of Table 4, show that all strategy types exhibit improved predictive outcomes (i.e., stronger positive return predictive power) when associated with longer predictive horizons. Specifically, messages emphasizing daily horizon negatively forecast subsequent returns across all three strategy categories (TA, FA, OS). In particular, long-term FA sentiment significantly and positively predicts returns over the window $t + 6$ to $t + 10$. These findings further suggest that the positive informativeness of FA sentiment primarily derives from discussions focused on long-term investment horizons, whereas the negative predictive power associated with TA and OS sentiments is largely driven by investors' excessive attention to short-term trading horizons.

4.3. Investor Experience

Investment performance may vary significantly with investor experience. The literature documents numerous technical indicators capable of generating superior returns (e.g., Jegadeesh and Titman, 1993b; Han et al., 2013, 2016; Jiang et al.,

[19]The fraction of messages classified with identifiable predictive horizons is comparable to those classified by strategy type. We observe considerable overlap between strategy and horizon-classified message groups, likely because descriptions of investment strategies often implicitly include information about intended holding periods.

2023; Murray et al., 2024). It is therefore plausible that sophisticated investors on the StockTwits platform can better identify valuable technical signals, achieving higher investment returns relative to the average TA user.

To investigate this possibility, we interact each sentiment measure with the fraction of messages contributed by self-reported professional investors for each stock-day. The results of this analysis are presented in Table 5. We find strong evidence that greater participation from professional users significantly enhances the predictive power of TA-related sentiment. Similarly, higher professional participation is positively associated with improved predictive performance for OS sentiment. However, we find little effect of professionals on the predictive ability of FA sentiment. This result may reflect a narrower dispersion in fundamental analysis skills among investors on the platform, suggesting that self-declared professional users on Stocktwits are not critical in identifying profitable fundamental signals.

4.4. Event Study: GME Short Squeeze

Bradley et al. (2024) provide empirical evidence that the GameStop (GME) short squeeze episode in early 2021 significantly reduced the informativeness of investor-generated content on social media platforms such as WallStreetBets. Motivated by their findings, we further investigate whether the impact of this event varies across distinct retail investors' different strategies. To conduct this event-study analysis, we focus on the period spanning 180 days before and after January 12, 2021—the starting point of the GameStop short squeeze. We interact our sentiment measures for each strategy type (TA, FA, OS, NS) with an indicator denoting the post-GME event period (January 14 onwards).

Our results, reported in Table 6, indicate a general decline in the informativeness (i.e., positive return predictability) of sentiments across all investment strategies following the GME event. However, this reduction is most pronounced within the TA strategy category. The disproportionate decline in TA informativeness likely stems from the substantial influx of inexperienced retail investors

drawn to the market by the high-profile nature of the short squeeze. Our previous finding with investor experience supports this hypothesis: professional investors typically produce more informative and reliable TA discussions. Therefore, we posit that the surge of novice participants disproportionately gravitated toward TA strategies—perhaps due to their greater accessibility and perceived simplicity—leading to a dilution in the quality and predictive value of TA-based sentiment. Our findings draw broader implications of market disruptions such as the GME short squeeze, highlighting how shifts in the composition and experience levels of market participants, along with their heterogeneous preferences in the usage of various strategies, can substantially influence the effectiveness and reliability of investor-generated social media discourse.

5. Retail Strategies and Retail Trading

5.1. Retail Market Order Flows

Our analysis thus far shows that StockTwits sentiments derived from different investment strategies yield substantially different investment outcomes. However, it remains possible that investors may "talk the talk" but not "walk the talk," implying that retail investor discussions on social media might not accurately reflect their actual trading behavior. Since StockTwits does not provide users' brokerage trading data, we cannot directly test whether an investor's trades align with their posted messages. In this section, we attempt to infer the representativeness of StockTwits sentiment in capturing retail investor beliefs and trading by examining its contemporaneous relationship with aggregate retail market order flows.

We identify retail market orders and compute two alternative measures of retail market order imbalance (*OIB*): $OIB_{i,t}^{BJZZ}$, following the algorithm in Boehmer et al. (2021), and $OIB_{i,t}^{BHJOS}$, based on the modified methodology of Barber et al. (2022).

We estimate the following panel regression at the stock-day level:

$$OIB_{i,t} = \sum_{type} \beta_1^{type} Sentiment_{i,t}^{type,IN} + \sum_{type} \beta_2^{type} Sentiment_{i,t}^{type,OV}$$
$$+ \beta_3 Attention_{i,t} + \gamma \boldsymbol{X_{i,t}} + \delta_t + \epsilon_{i,t}.$$

To align closely with market orders, in this test we separately calculate sentiment measures using messages posted during intraday (IN) trading hours (i.e., 9:30-16:00) and overnight (OV) intervals (from the prior market close to the market open).

The regression results are presented in Panel A of Table 7. Columns (1)-(3) focus on the retail market order imbalance from Boehmer et al. (2021) (OIB^{BJZZ}). Column (1) demonstrates that intraday sentiments of all four strategies exhibit significant and positive relationships with retail market order imbalance. Column (2) examines overnight sentiment and finds positive yet substantially weaker associations with the intraday retail order imbalance. Column (3), which includes both intraday and overnight sentiments simultaneously, indicates that intraday sentiment maintains significant explanatory power, while the overnight sentiment measures generally become insignificant.

Columns (4)-(6) repeat the regression analysis using the alternative retail market order imbalance from Barber et al. (2023) (OIB^{BHJOS}). The results are qualitatively similar: intraday sentiments consistently demonstrate much stronger relationships with retail order imbalance than their overnight counterparts, across all strategy categories.

Overall, our findings provide robust evidence that intraday StockTwits sentiment indeed effectively captures contemporaneous retail investor trading during market hours.

5.2. Retail Strategies and Retail Order Flow Informativeness

Boehmer et al. (2021) demonstrate that aggregate retail market order imbal-

ances are informative and positively predict future stock returns. Building on our earlier finding that StockTwits sentiment closely aligns with retail investor beliefs and trading strategies, we now investigate how different investor strategy types influence the informativeness of retail order flows.

We first perform a pooled regression of *OIB* on the four intraday sentiment measures categorized by investment strategy type. We decompose the aggregate *OIB* into strategy-specific components as follows:

$$OIB_{i,t}^{type} = \beta_1^{type} \times Sentiment_{i,t}^{type,IN}.$$

We interpret the residual from this decomposition as capturing retail order flow that cannot be explained by retail investors' publicly shared analyses on StockTwits. This residual component likely includes either noise trading or private information held by retail investors but not publicly expressed through social media.

We then regress next-day stock returns on each of these decomposed components of *OIB*, including the residual term:

$$Return_{i,t} = \beta_1 OIB_{i,t}^{TA} + \beta_2 OIB_{i,t}^{FA} + \beta_3 OIB_{i,t}^{OS} + \beta_4 OIB_{i,t}^{NS} + \beta_5 OIB_{i,t}^{Resid} + \gamma X_{i,t} + \delta_t + \epsilon_{i,t+1}.$$

Results are reported in Panel B of Table 7. Column (1) focuses on the decomposition based on the OIB^{BJZZ} measure, while column (2) employs the alternative OIB^{BHJOS} measure. Several key findings emerge. First, TA- and OS-based components of *OIB* negatively predict future stock returns, indicating that these strategies reduce the informativeness of aggregate retail market orders. In contrast, FA-based *OIB* positively predicts returns, thereby enhancing informativeness. Retail order imbalance attributable to non-strategy (NS) messages exhibits minimal predictive power. Notably, the residual component positively and significantly predicts stock returns, suggesting that retail order flow includes private information distinct from strategies inferred from StockTwits messages.

5.3. *Retail Strategies and Robinhood Herding*

An important aspect of retail investor behavior is "attention-driven trading," which is especially prevalent among new investors on zero-commission trading platforms such as Robinhood (Barber et al., 2022). In this subsection, we investigate whether investor sentiments derived from different strategy types are related to retail investor herding activities observed on Robinhood.

Following Barber et al. (2022) and Welch (2022), we process RobinTrack data to identify Robinhood buy-herding events ($RH\ Herding_{i,t}$). Specifically, we create an indicator variable equal to one if stock i ranks among the top 10 stocks on day t based on the daily percentage increase in Robinhood users holding the stock, provided that at least 100 Robinhood users held the stock at the end of day $t-1$. Our sample spans May 2018 to August 2020, aligning with the availability of Robinhood user account data from RobinTrack.

To assess the relationship between StockTwits sentiment and Robinhood herding events, we estimate the following stock-day panel regression:

$$RH\ Herding_{i,t} = \sum_{type} \beta_1^{type} Sentiment_{i,t}^{type} + \beta_2 Attention_{i,t} + \gamma \boldsymbol{X_{i,t}} + \delta_t + \eta_i + \epsilon_{i,t}.$$

Table 8 presents the regression results. Columns (1) through (4) separately analyze the predictive relationship between Robinhood herding and StockTwits sentiment for each investment strategy category (TA, FA, OS, NS). We find that sentiment from each strategy category is positively associated with Robinhood herding events. However, the economic magnitude of this association is notably stronger for TA sentiment compared to other strategy types. Column (5) simultaneously includes all sentiment categories, and the results remain consistent: sentiment across all types continues to positively predict Robinhood herding, with TA sentiment demonstrating the strongest relationship.

Overall, our findings indicate that StockTwits sentiment effectively captures retail herding behavior on Robinhood, and that the strength of this relationship

varies substantially by investment strategy type. In particular, our evidence suggests that technical analysis strategies are especially prone to crowded trading, potentially explaining the relatively poorer performance outcomes associated with these strategies.

6. Conclusion

In conclusion, this study demonstrates the potential of combining rich social media data and large language models (LLMs) to better understand investors trading strategies, sentiment, trading, and market outcomes. We use LLMs to identify strategies embedStockTwits messages categorized by trading strategies—technical analysis (TA), fundamental analysis (FA), other strategies (OS), and messages without a clearly defined strategy (NS). Our analysis uncovers that retail investors dynamically adapt their trading approaches based on the prevailing informational environment. Specifically, we observe that investors increasingly rely on TA during periods with limited firm-specific news, while FA becomes more prominent as news flow intensifies. Our empirical findings reveal a distinct relationship between message sentiment and stock returns, with TA sentiments negatively predicting returns, particularly among short-term and less experienced traders, whereas FA sentiments positively predict returns. Additionally, our results show that all message categories positively correlate with retail order imbalance, underscoring the significant role of social media-derived retail sentiment in influencing actual trading behaviors. Notably, TA-driven sentiment is most closely associated with retail herding behaviors, often resulting in subsequent price reversals. Overall, this research highlights the substantial potential of leveraging social media data and advanced language modeling techniques to gain deeper insights into investor decision-making processes.

REFERENCES

Abarbanell, Jeffrey S, and Brian J Bushee, 1997, Fundamental analysis, future earnings, and stock prices, *Journal of accounting research* 35, 1–24.

Antweiler, Werner, and Murray Z Frank, 2004, Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards, *The Journal of Finance* 59, 1259–1294.

Barber, Brad M, Xing Huang, Philippe Jorion, Terrance Odean, and Christopher Schwarz, 2023, A (Sub)penny for Your Thoughts: Tracking Retail Investor Activity in TAQ, *The Journal of Finance* .

Barber, Brad M, Xing Huang, Terrance Odean, and Christopher Schwarz, 2022, Attention-Induced Trading and Returns: Evidence from Robinhood Users, *The Journal of Finance* 77, 3141–3190.

Barber, Brad M, and Terrance Odean, 2000, Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors, *The journal of Finance* 55, 773–806.

Barber, Brad M, and Terrance Odean, 2008, All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors, *The Review of Financial Studies* 21, 785–818.

Barber, Brad M, and Terrance Odean, 2013, The behavior of individual investors, in *Handbook of the Economics of Finance*, volume 2, 1533–1570 (Elsevier).

Blume, Lawrence, David Easley, and Maureen O'hara, 1994, Market Statistics and Technical Analysis: The Role of Volume, *The Journal of Finance* 49, 153–181.

Boehmer, Ekkehart, Charles M Jones, Xiaoyan Zhang, and Xinran Zhang, 2021, Tracking Retail Investor Activity, *The Journal of Finance* 76, 2249–2305.

Bradley, Daniel, Jan Hanousek Jr, Russell Jame, and Zicheng Xiao, 2024, Place Your Bets? The Value of Investment Research on Reddit's Wallstreetbets, *The Review of Financial Studies* 37, 1409–1459.

Brock, William, Josef Lakonishok, and Blake LeBaron, 1992, Simple Technical Trading Rules and The Stochastic Properties of Stock Returns, *The Journal of Finance* 47, 1731–1764.

Brown, David P, and Robert H Jennings, 1989, On Technical Analysis, *The Review of Financial Studies* 2, 527–551.

Bryzgalova, Svetlana, Anna Pavlova, and Taisiya Sikorskaya, 2023, Retail Trading in Options and the Rise of the Big Three Wholesalers, *The Journal of Finance* 78, 3465–3514.

Chen, Hailiang, Prabuddha De, Yu Hu, and Byoung-Hyoun Hwang, 2014, Wisdom of Crowds: The Value of Stock Opinions Transmitted Through Social Media, *The Review of Financial Studies* 27, 1367–1403.

Chen, Hailiang, and Byoung-Hyoun Hwang, 2022, Listening in on Investors' Thoughts and Conversations, *Journal of Financial Economics* 145, 426–444.

Chinco, Alex, Samuel M Hartzmark, and Abigail B Sussman, 2022, A New Test of Risk Factor Relevance, *The Journal of Finance* 77, 2183–2238.

Choi, James J, and Adriana Z Robertson, 2020, What Matters to Individual Investors? Evidence from The Horse's Mouth, *The Journal of Finance* 75, 1965–2020.

Cookson, J Anthony, Chukwuma Dim, and Marina Niessner, 2024a, Disagreement on the Horizon, *Available at SSRN* .

Cookson, J Anthony, Joseph E Engelberg, and William Mullins, 2023, Echo Chambers, *The Review of Financial Studies* 36, 450–500.

Cookson, J Anthony, Runjing Lu, William Mullins, and Marina Niessner, 2024b, The Social Signal, *Journal of Financial Economics* 158, 103870.

Cookson, J Anthony, William Mullins, and Marina Niessner, 2024c, Social Media and Finance, *Available at SSRN 4806692* .

Cookson, J Anthony, and Marina Niessner, 2020, Why Don't We Agree? Evidence from A Social Network of Investors, *The Journal of Finance* 75, 173–228.

Coval, Joshua D, David Hirshleifer, and Tyler Shumway, 2021, Can Individual Investors Beat The Market?, *The Review of Asset Pricing Studies* 11, 552–579.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, 2018, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, *CoRR* abs/1810.04805.

Eaton, Gregory W, T Clifton Green, Brian S Roseman, and Yanbin Wu, 2022, Retail Trader Sophistication and Stock Market Quality: Evidence from Brokerage Outages, *Journal of Financial Economics* 146, 502–528.

Eisfeldt, Andrea L, Gregor Schubert, and Miao Ben Zhang, 2023, Generative AI and Firm Values, Technical report, National Bureau of Economic Research.

Farrell, Michael, T Clifton Green, Russell Jame, and Stanimir Markov, 2022, The Democratization of Investment Research and The Informativeness of Retail Investor Trading, *Journal of Financial Economics* 145, 616–641.

George, Thomas J, and Chuan-Yang Hwang, 2004, The 52-Week High and Momentum Investing, *The Journal of Finance* 59, 2145–2176.

Giannini, Robert, Paul Irvine, and Tao Shu, 2018, Nonlocal Disadvantage: An Examination of Social Media Sentiment, *The Review of Asset Pricing Studies* 8, 293–336.

Giannini, Robert, Paul Irvine, and Tao Shu, 2019, The Convergence and Divergence of Investors' Opinions Around Earnings News: Evidence from A Social Network, *Journal of Financial Markets* 42, 94–120.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2021, Five Facts about Beliefs and Portfolios, *American Economic Review* 111, 1481–1522.

González-Carvajal, Santiago, and Eduardo C Garrido-Merchán, 2020, Comparing BERT against Traditional Machine Learning Text Classification, *arXiv preprint arXiv:2005.13012* .

Gu, Yuxian, Li Dong, Furu Wei, and Minlie Huang, 2023, Knowledge Distillation of Large Language Models, *arXiv preprint arXiv:2306.08543* .

Han, Yufeng, Ke Yang, and Guofu Zhou, 2013, A New Anomaly: The Cross-Sectional Profitability of Technical Analysis, *Journal of Financial and Quantitative Analysis* 48, 1433–1461.

Han, Yufeng, Guofu Zhou, and Yingzi Zhu, 2016, A Trend Factor: Any Economic Gains from Using Information over Investment Horizons?, *Journal of Financial Economics* 122, 352–375.

Hinton, G, 2015, Distilling the Knowledge in A Neural Network, *arXiv preprint arXiv:1503.02531* .

Hirshleifer, David, Lin Peng, and Qiguang Wang, Forthcoming, News Diffusion In Social Networks and Stock Market Reactions, *The Review of Financial Studies* .

Huang, Linmei, Bill Qiao, and Dexin Zhou, 2024, Are Memes a Sideshow: Evidence from WallStreetBets, *Available at SSRN 4785481* .

Jegadeesh, Narasimhan, 1991, Seasonality in Stock Price Mean Reversion: Evidence from The US and The UK, *The Journal of Finance* 46, 1427–1444.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993a, Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency, *The Journal of Finance* 48, 65–91.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993b, Returns to buying winners and selling losers: Implications for stock market efficiency, *The Journal of finance* 48, 65–91.

Jha, Manish, Jialin Qian, Michael Weber, and Baozhong Yang, 2024, ChatGPT and Corporate Policies, Technical report, National Bureau of Economic Research.

Jiang, Jingwen, Bryan Kelly, and Dacheng Xiu, 2023, (re-) Imag (in) ing Price Trends, *The Journal of Finance* 78, 3193–3249.

Jiang, Jingwen, Bryan T Kelly, and Dacheng Xiu, 2024a, Expected Returns and Large Language Models, *Available at SSRN* .

Jiang, Zhengyang, Cameron Peng, and Hongjun Yan, 2024b, Personality Differences and Investment Decision-Making, *Journal of Financial Economics* 153, 103776.

Kakhbod, Ali, Seyed Mohammad Kazempour, Dmitry Livdan, and Norman Schuerhoff, 2023, Finfluencers, *Available at SSRN 4428232* .

Kaniel, Ron, Gideon Saar, and Sheridan Titman, 2008, Individual Investor Trading and Stock Returns, *The Journal of Finance* 63, 273–310.

Kelley, Eric K, and Paul C Tetlock, 2013, How Wise Are Crowds? Insights from Retail Orders and Stock Returns, *The Journal of Finance* 68, 1229–1265.

Korinek, Anton, 2023, Generative AI for Economic Research: Use Cases and Implications for Economists, *Journal of Economic Literature* 61, 1281–1317.

Kumar, Alok, and Charles MC Lee, 2006, Retail Investor Sentiment and Return Comovements, *The Journal of Finance* 61, 2451–2486.

Laudenbach, Christine, Annika Weber, Rüdiger Weber, and Johannes Wohlfart, 2024, Beliefs About the Stock Market and Investment Choices: Evidence from a Survey and a Field Experiment, *The Review of Financial Studies* Forthcoming.

Li, Kai, Feng Mai, Rui Shen, Chelsea Yang, and Tengfei Zhang, 2023, Dissecting Corporate Culture Using Generative AI–Insights from Analyst Reports, *Available at SSRN 4558295* .

Liu, Hongqi, Cameron Peng, Wei A Xiong, and Wei Xiong, 2022, Taming The Bias Zoo, *Journal of Financial Economics* 143, 716–741.

Lo, Andrew W, Harry Mamaysky, and Jiang Wang, 2000, Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation, *The Journal of Finance* 55, 1705–1765.

Lopez-Lira, Alejandro, and Yuehua Tang, 2023, Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models, *arXiv preprint arXiv:2304.07619* .

Murray, Scott, Yusen Xia, and Houping Xiao, 2024, Charting by Machines, *Journal of Financial Economics* 153, 103791.

Porta, Rafael La, Josef Lakonishok, Andrei Shleifer, and Robert Vishny, 1997, Good news for value stocks: Further evidence on market efficiency, *the Journal of Finance* 52, 859–874.

Welch, Ivo, 2022, The Wisdom of the Robinhood Crowd, *The Journal of Finance* 77, 1489–1527.

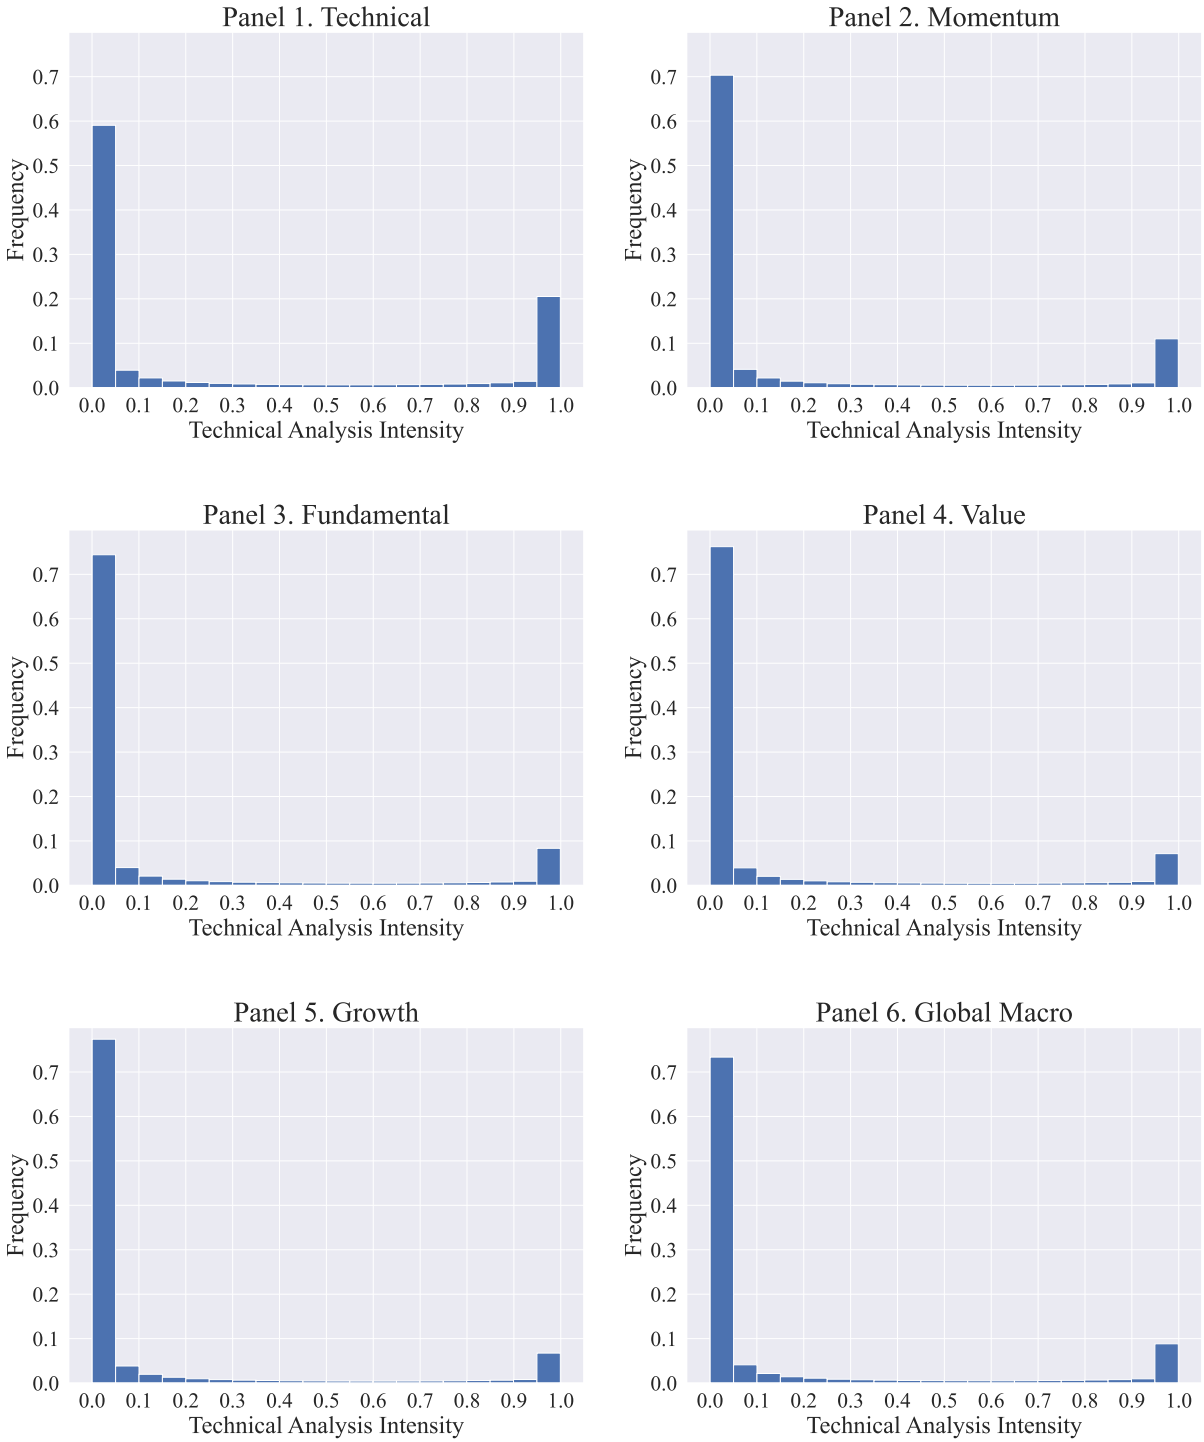


Fig. 1. **Distribution of Technical Analysis Intensity at the Message Level by Self-Declared Investment Approach**

This figure shows the distribution of technical analysis intensity scores for messages posted on Stocktwits, grouped by users' self-declared investment approaches. Each message is assigned a probabilistic technical analysis score ("Technical Analysis Intensity") generated by our fine-tuned BERT model. Higher intensity indicates a greater likelihood that the message employs technical analysis. The histograms illustrate the frequency of messages across different intensity levels for each investor approach subgroup.

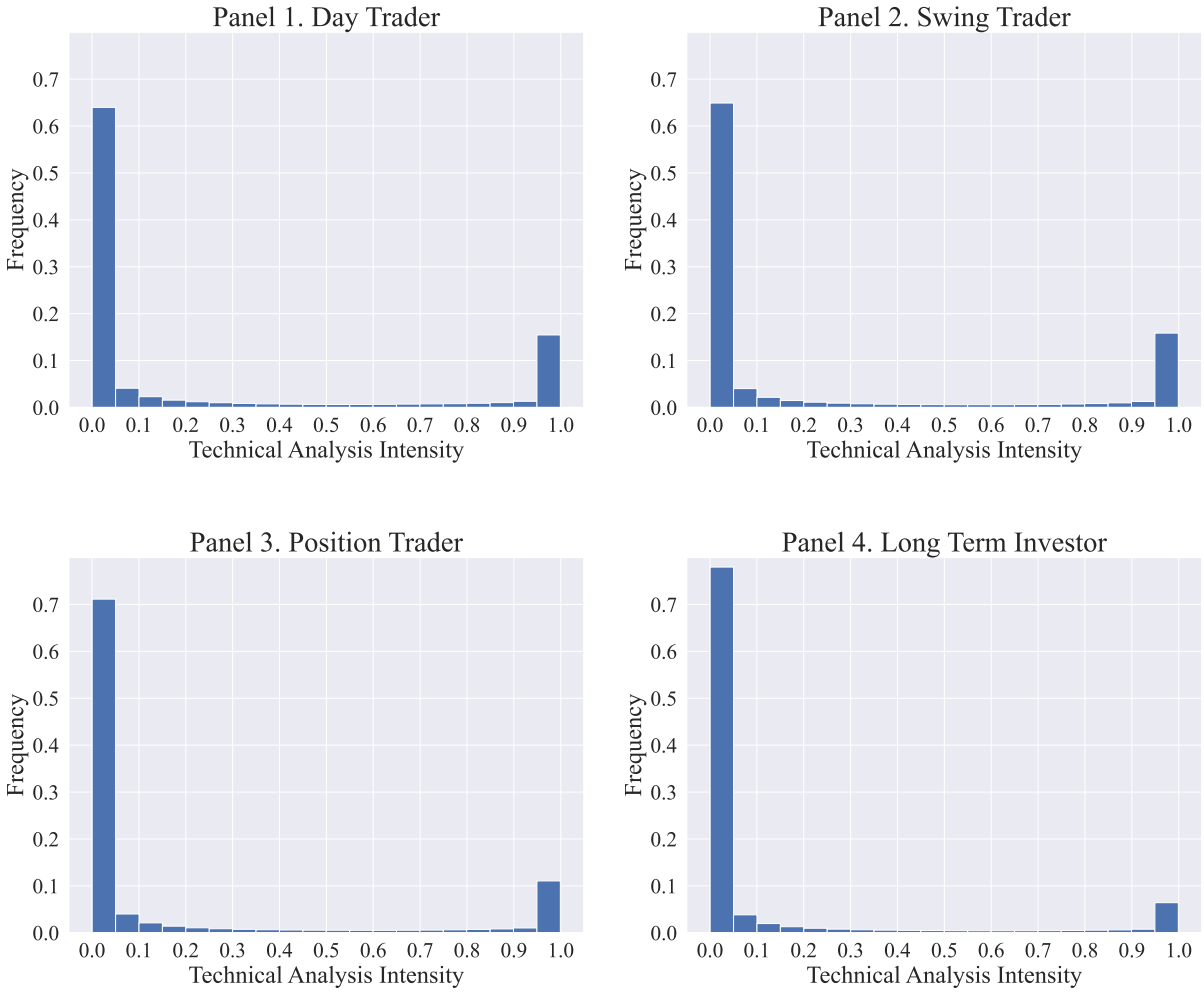


Fig. 2. **Distribution of Technical Analysis Intensity at the Message Level by Self-Declared Investment Horizon**

This figure shows the distribution of technical analysis intensity scores for messages posted on Stocktwits, grouped by users' self-declared investment horizons. Each message is assigned a probabilistic technical analysis score ("Technical Analysis Intensity") generated by our fine-tuned BERT model. Higher intensity indicates a greater likelihood that the message employs technical analysis. The histograms illustrate the frequency of messages across different intensity levels for each investor horizon subgroup.

Panel A. Word Cloud of Unigrams in Technical Messages



Panel B. Word Cloud of Bigrams in Technical Messages

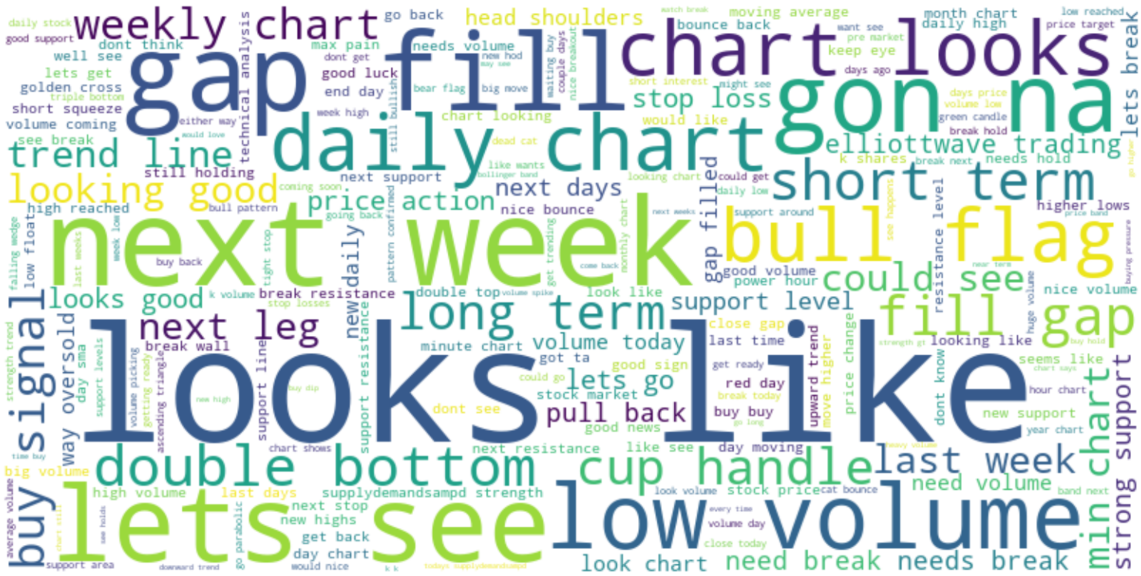


Fig. 3. **Word Clouds of Technical Analysis Messages**

This figure presents word clouds derived from StockTwits messages classified as technical messages (Technical Analysis Intensity $\geq$ 0.95). Panel A displays the word cloud for the most frequent unigram (single-word) terms, and Panel B shows the word cloud for the most frequent bigram (two-word) phrases.

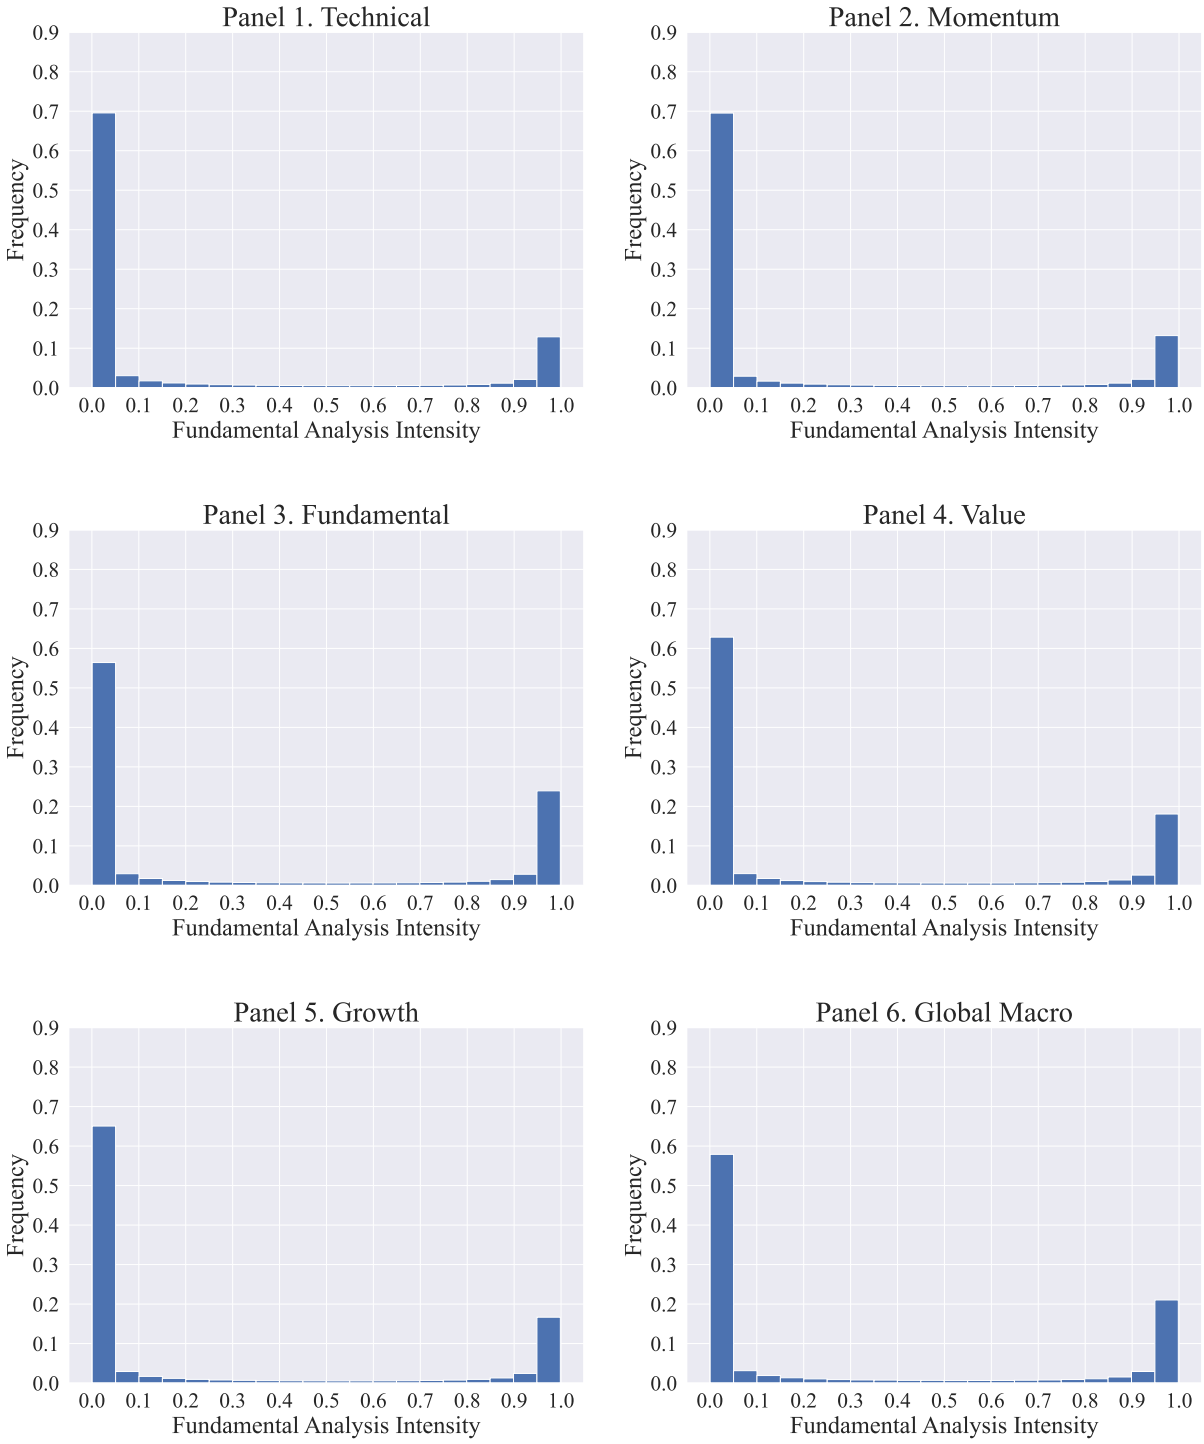


Fig. 4. **Distribution of Fundamental Analysis Intensity at the Message Level by Self-Declared Investment Approach**

This figure shows the distribution of fundamental analysis intensity scores for messages posted on Stocktwits, grouped by users' self-declared investment approaches. Each message is assigned a probabilistic fundamental analysis score ("Fundamental Analysis Intensity") generated by our fine-tuned BERT model. Higher intensity indicates a greater likelihood that the message employs fundamental analysis. The histograms illustrate the frequency of messages across different intensity levels for each investor approach subgroup.

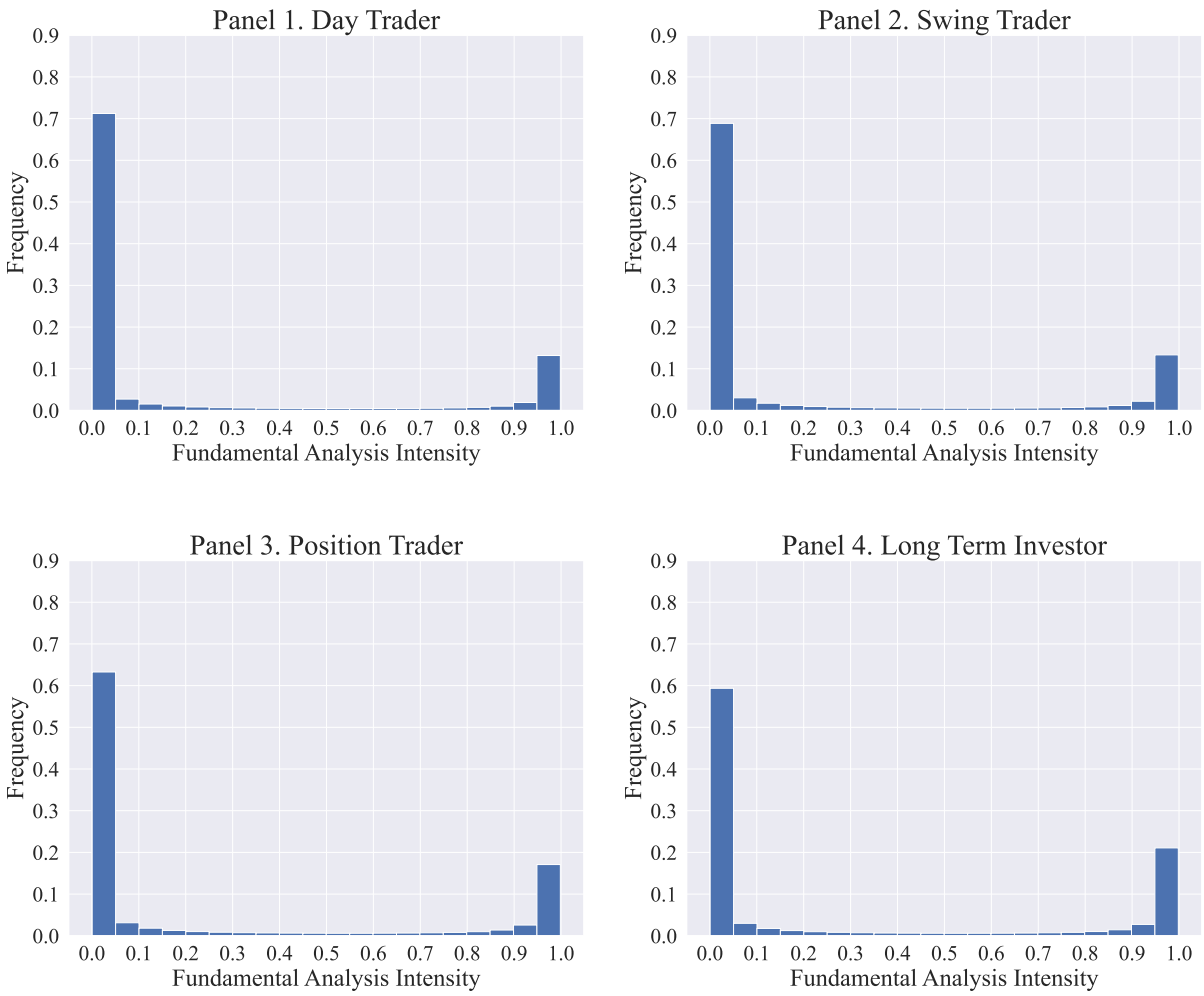


Fig. 5. **Distribution of Fundamental Analysis Intensity at the Message Level by Self-Declared Investment Horizon**

This figure shows the distribution of fundamental analysis intensity scores for messages posted on Stocktwits, grouped by users' self-declared investment horizons. Each message is assigned a probabilistic fundamental analysis score ("Fundamental Analysis Intensity") generated by our fine-tuned BERT model. Higher intensity indicates a greater likelihood that the message employs fundamental analysis. The histograms illustrate the frequency of messages across different intensity levels for each investor horizon subgroup.

Panel A. Word Cloud of Unigrams in Fundamental Messages



Panel B. Word Cloud of Bigrams in Fundamental Messages



Fig. 6. **Word Clouds of Fundamental Analysis Messages**

This figure presents word clouds derived from StockTwits messages classified as fundamental messages (Fundamental Analysis Intensity ≥ 0.95). Panel A displays the word cloud for the most frequent unigram (single-word) terms, and Panel B shows the word cloud for the most frequent bigram (two-word) phrases.

Table 1: **Summary Statistics**

Panel A reports summary statistics for the message-level sample, restricted to messages for which users' self-reported biographic characteristics are available. The reported variables include the frequency of identified trading strategies (Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS)), self-reported investor profiles (Technical Investor, Long-Term Investor, Swing/Day Trader, Professional, and Novice), message length (number of words), and TF-IDF weighted keyword scores for technical and fundamental analysis based on word lists from Cookson and Niessner (2020). Panel B presents summary statistics for variables in the stock-day level sample, including sentiment scores derived from StockTwits messages categorized by investment strategies: Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), and Non-Strategy (NS). We also report investor attention, retail order imbalance (OIB) measures based on methodologies from Boehmer et al. (2021) and Barber et al. (2023), as well as firm characteristics including market capitalization, book-to-market ratio, asset growth, gross profitability, analyst coverage, institutional ownership, maximum daily return in the prior month (MAX), abnormal turnover, and abnormal news article volume. The sample period spans January 2010 to June 2023. Detailed variable definitions are provided in Appendix Table A.2.

Panel A: Message-Level Sample with Self-reported User Information

	N	Mean	Median	StdDev	10th	25th	75th	90th
Usage^{TA}	21,641,362	0.14	0.00	0.35	0.00	0.00	0.00	1.00
Usage^{FA}	21,641,362	0.17	0.00	0.37	0.00	0.00	0.00	1.00
Usage^{OS}	21,641,362	0.11	0.00	0.31	0.00	0.00	0.00	1.00
Technical Investor	21,641,362	0.57	1.00	0.50	0.00	0.00	1.00	1.00
Long-Term Investor	21,641,362	0.22	0.00	0.42	0.00	0.00	0.00	1.00
Swing or Day Trader	21,641,362	0.59	1.00	0.49	0.00	0.00	1.00	1.00
Professional	21,641,362	0.32	0.00	0.46	0.00	0.00	1.00	1.00
Novice	21,641,362	0.14	0.00	0.35	0.00	0.00	0.00	1.00
Number of Words	21,641,362	14.70	10.00	17.19	3.00	5.00	18.00	27.00
Technical^{TF-IDF}	21,641,362	0.02	0.00	0.06	0.00	0.00	0.01	0.05
Fundamental^{TF-IDF}	21,641,362	0.01	0.00	0.03	0.00	0.00	0.00	0.02

Panel B: Stock-Level Sample

	N	Mean	Median	StdDev	10th	25th	75th	90th
Sentiment^{TA}	2,974,934	0.23	0.00	0.53	0.00	0.00	1.00	1.00
Sentiment^{FA}	2,974,934	0.29	0.00	0.57	0.00	0.00	1.00	1.00
Sentiment^{OS}	2,974,934	0.16	0.00	0.50	0.00	0.00	0.25	1.00
Sentiment^{NS}	2,974,934	0.26	0.00	0.62	-1.00	0.00	1.00	1.00
Attention	2,974,934	0.09	0.01	0.62	0.00	0.00	0.04	0.15
OIB^{BJZZ}	2,974,934	-0.01	-0.00	0.27	-0.33	-0.14	0.12	0.29
OIB^{BHJOS}	2,974,934	-0.01	0.00	0.25	-0.28	-0.11	0.11	0.25
Log(Market Cap)	2,974,934	7.12	7.14	2.52	3.68	5.25	8.95	10.53
Book-to-Market	2,974,934	0.63	0.39	0.82	0.09	0.19	0.77	1.33
Asset Growth	2,974,934	1.07	1.00	0.54	0.89	0.96	1.05	1.17
Gross Profit-to-Asset	2,974,934	0.05	0.05	0.09	-0.04	0.01	0.09	0.15
Number of Analysts	2,974,934	9.74	7.00	8.98	1.00	3.00	15.00	23.00
Institutional Ownership	2,974,934	0.60	0.69	0.32	0.09	0.33	0.86	0.96
MAX	2,974,934	0.08	0.06	0.09	0.02	0.03	0.10	0.16
Abnormal Turnover	2,974,934	-0.11	-0.11	0.63	-0.81	-0.44	0.22	0.61
Abnormal News	2,974,934	-0.50	-0.53	0.98	-1.71	-1.15	0.00	0.65

Table 2: **Determinants of Retail Investor Strategy Usage Identified via Large Language Model**

This table reports panel regression results analyzing determinants of retail investors' usage of different trading strategies, classified from StockTwits messages using a large language model. Panels A, B, and C correspond to Technical Analysis (TA), Fundamental Analysis (FA), and Other Strategies (OS), respectively. The regression is estimated at the individual message level, following the specification:

$$Usage_{i,j,t,n}^{type} = \beta_1 X_j^{investor} + \beta_2 Y_{i,t}^{stock} + \beta_3 Z_{j,i,t,n}^{message} + FEs + \epsilon_{i,j,t,n},$$

where $Usage_{i,j,t,n}^{type}$ is an indicator variable equal to one if message n, posted by investor j about stock i on day t, is classified into strategy type TA, FA, or OS by our BERT models. Investor-level characteristics ($X_j^{investor}$) include self-reported indicators such as Technical Investor, Long-Term Investor, and Professional Investor. Stock-day characteristics ($Y_{i,t}^{stock}$) include abnormal turnover and abnormal news article volume. Message-specific attributes ($Z_{j,i,t,n}^{message}$) include the log of the number of words and TF-IDF weighted scores for technical and fundamental keywords. Column (4) includes trading-date and stock fixed effects. Column (5) further incorporates investor fixed effects. Column (6) includes date fixed effects along with stock-investor interacted fixed effects. Standard errors are clustered by investor, stock, and day, with t-statistics reported in brackets. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Determinants of Technical Analysis Usage

	(1)	(2)	(3)	(4)	(5)	(6)
Technical Investor$_j$	0.103***	0.075***	0.076***	0.068***		
	[19.74]	[17.50]	[17.60]	[17.42]		
Swing or Day Trader$_j$		0.020***	0.022***	0.022***		
		[3.39]	[3.89]	[4.22]		
Long-Term Investor$_j$		-0.028***	-0.028***	-0.023***		
		[-6.10]	[-6.44]	[-5.65]		
Professional$_j$		0.047***	0.044***	0.033***		
		[6.45]	[6.23]	[5.30]		
Novice$_j$		-0.022***	-0.018***	-0.013***		
		[-7.46]	[-6.21]	[-5.31]		
Abnormal Turnover$_{i,t}$			-0.003***	-0.005***	-0.003***	-0.003***
			[-2.99]	[-6.37]	[-7.45]	[-8.10]
Abnormal News$_{i,t}$			-0.006***	-0.005***	-0.003***	-0.002***
			[-9.73]	[-11.86]	[-12.15]	[-10.15]
Technical$_{i,j,t,n}^{TF-IDF}$			0.825***	0.798***	0.683***	0.653***
			[40.95]	[40.26]	[42.09]	[39.56]
Fundamental$_{i,j,t,n}^{TF-IDF}$			-0.664***	-0.632***	-0.533***	-0.498***
			[-27.39]	[-26.73]	[-28.29]	[-28.52]
Log(Number of Words$_{i,j,t,n}$)			0.047***	0.049***	0.044***	0.043***
			[16.62]	[18.92]	[28.62]	[26.76]
Date FEs	No	No	No	Yes	Yes	Yes
Stock FEs	No	No	No	Yes	Yes	No
Investor FEs	No	No	No	No	Yes	No
Stock × Investor FEs	No	No	No	No	No	Yes
N	21,641,362	21,641,362	21,641,362	21,641,218	21,623,813	20,630,883
R^2	0.022	0.030	0.063	0.084	0.213	0.287

Panel B. Determinants of Fundamental Analysis Usage

	(1)	(2)	(3)	(4)	(5)	(6)
Technical Investor$_j$	-0.073***	-0.052***	-0.041***	-0.042***		
	[-12.02]	[-11.12]	[-9.90]	[-11.09]		
Swing or Day Trader$_j$		-0.026***	-0.016***	-0.015***		
		[-5.16]	[-3.40]	[-3.54]		
Long-Term Investor$_j$		0.038***	0.033***	0.032***		
		[4.80]	[4.40]	[4.77]		
Professional$_j$		0.034***	0.029***	0.018***		
		[4.06]	[3.97]	[3.06]		
Novice$_j$		-0.031***	-0.020***	-0.016***		
		[-8.80]	[-6.55]	[-5.90]		
Abnormal Turnover$_{i,t}$			-0.005***	-0.009***	-0.007***	-0.006***
			[-5.51]	[-13.45]	[-16.46]	[-14.35]
Abnormal News$_{i,t}$			0.004***	0.006***	0.005***	0.005***
			[4.72]	[11.17]	[13.80]	[13.76]
Technical$^{TF-IDF}_{i,j,t,n}$			-0.313***	-0.320***	-0.256***	-0.234***
			[-20.29]	[-22.38]	[-24.32]	[-23.63]
Fundamental$^{TF-IDF}_{i,j,t,n}$			0.768***	0.760***	0.639***	0.567***
			[15.74]	[16.87]	[18.32]	[17.12]
Log(Number of Words$_{i,j,t,n}$)			0.138***	0.137***	0.127***	0.125***
			[41.01]	[47.24]	[63.95]	[58.25]
Date FEs	No	No	No	Yes	Yes	Yes
Stock FEs	No	No	No	Yes	Yes	No
Investor FEs	No	No	No	No	Yes	No
Stock × Investor FEs	No	No	No	No	No	Yes
N	21,641,362	21,641,362	21,641,362	21,641,218	21,623,813	20,630,883
R^2	0.009	0.016	0.140	0.161	0.260	0.325

Panel C. Determinants of Other Strategy Usage

	(1)	(2)	(3)	(4)	(5)	(6)
Technical Investor$_j$	0.032***	0.017***	0.018***	0.014***		
	[14.43]	[8.53]	[8.98]	[7.71]		
Swing or Day Trader$_j$		0.011***	0.013***	0.012***		
		[5.01]	[5.35]	[5.23]		
Long-Term Investor$_j$		-0.016***	-0.016***	-0.013***		
		[-6.94]	[-7.16]	[-5.90]		
Professional$_j$		0.022***	0.020***	0.015***		
		[7.44]	[6.79]	[5.70]		
Novice$_j$		-0.015***	-0.013***	-0.011***		
		[-9.78]	[-8.46]	[-7.43]		
Abnormal Turnover$_{i,t}$			0.005***	0.004***	0.002***	0.002***
			[10.36]	[11.89]	[12.65]	[9.32]
Abnormal News$_{i,t}$			-0.001***	-0.001***	-0.001***	-0.001***
			[-5.14]	[-6.44]	[-7.01]	[-7.38]
Technical$^{TF-IDF}_{i,j,t,n}$			0.228***	0.220***	0.191***	0.185***
			[16.49]	[15.55]	[15.48]	[15.10]
Fundamental$^{TF-IDF}_{i,j,t,n}$			-0.289***	-0.275***	-0.231***	-0.211***
			[-23.19]	[-23.03]	[-23.04]	[-20.94]
Log(Number of Words$_{i,j,t,n}$)			0.028***	0.029***	0.030***	0.032***
			[19.36]	[23.94]	[43.02]	[47.78]
Date FEs	No	No	No	Yes	Yes	Yes
Stock FEs	No	No	No	Yes	Yes	No
Investor FEs	No	No	No	No	Yes	No
Stock × Investor FEs	No	No	No	No	No	Yes
N	21,641,362	21,641,362	21,641,362	21,641,218	21,623,813	20,630,883
R^2	0.003	0.005	0.014	0.021	0.071	0.145

Table 3: **Predicting Next-Day Stock Returns Using Retail Investor Sentiment by Investment Strategy**

This table reports panel regression results examining the daily stock return predictability of retail investor sentiments by different strategy types. StockTwits messages are classified into four investment strategy categories using our BERT models: Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), and Non-Strategy messages (NS). We estimate predictive regressions at the stock-day level as follows:

$$Return_{i,t+1} = \beta_1 Sentiment_{i,t}^{type} + \beta_2 Attention_{i,t} + \gamma \boldsymbol{X_{i,t}} + \delta_t + \epsilon_{i,t+1},$$

where $Sentiment_{i,t}^{type}$ denotes investor sentiment toward stock i on day t, separately measured for each strategy type (TA, FA, OS, or NS). Following Cookson et al. (2024b), sentiment is computed as the difference between bullish and bearish messages normalized by their sum:

$$Sentiment_{i,t} = \frac{N_{i,t}^{Bullish} - N_{i,t}^{Bearish}}{N_{i,t}^{Bullish} + N_{i,t}^{Bearish}}.$$

$Attention_{i,t}$ captures retail investor attention to stock i on day t, defined as the percentage of total StockTwits messages that reference stock i on that day, $\frac{\#Messages_{i,t}}{\sum_i \#Messages_{i,t}} \times 100$. The vector $\boldsymbol{X_{i,t}}$ includes control variables: log market capitalization, log book-to-market ratio, asset growth, gross profitability, log analyst coverage, log institutional ownership, maximum daily return (MAX) in the previous month, abnormal turnover, abnormal news article volume, and lagged returns over the past five trading days. All regressions include trading day fixed effects. Standard errors are clustered by trading day, with t-statistics presented in brackets. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	Return$_{i,t+1}$ (%)				
	(1)	(2)	(3)	(4)	(5)
Sentiment$_{i,t}^{TA}$	-0.016**				-0.015**
	[-2.19]				[-2.25]
Sentiment$_{i,t}^{FA}$		0.014**			0.017***
		[2.17]			[2.88]
Sentiment$_{i,t}^{OS}$			-0.027***		-0.026***
			[-3.57]		[-3.73]
Sentiment$_{i,t}^{NS}$				-0.003	-0.002
				[-0.56]	[-0.30]
Attention$_{i,t}$	-0.056***	-0.056***	-0.056***	-0.056***	-0.056***
	[-5.47]	[-5.50]	[-5.46]	[-5.48]	[-5.47]
Stock Characteristics	Yes	Yes	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes	Yes	Yes
N	2,974,304	2,974,304	2,974,304	2,974,304	2,974,304
R^2	0.089	0.089	0.089	0.089	0.089

Table 4: **Predicting Stock Returns at Longer Horizons**

Panel A extends the analysis of daily return predictability from Table 3 to longer forecasting horizons (up to 15 days ahead). Specifically, we estimate the following panel regressions at the stock-day level:

$$Return_{i,t+h} = \beta_1 Sentiment_{i,t}^{type} + \beta_2 Attention_{i,t} + \gamma \boldsymbol{X_{i,t}} + \delta_t + \epsilon_{i,t+h},$$

where h denotes the return horizon: days $t+1$ to $t+5$, $t+6$ to $t+10$, or $t+11$ to $t+15$. Sentiments ($Sentiment_{i,t}^{type}$) are classified by our BERT models into Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), and Non-Strategy (NS).
Panel B explores whether the informativeness of StockTwits messages varies by the implied investment horizon (e.g., Long-term, Daily), which we also classify using BERT models. The panel regressions take the form:

$$Return_{i,t+h} = \beta_1 Sentiment_{i,t}^{type,horizon} + \beta_2 Attention_{i,t} + \gamma \boldsymbol{X_{i,t}} + \delta_t + \epsilon_{i,t+h},$$

where $Sentiment_{i,t}^{type,horizon}$ indicates investor sentiment toward stock i on day t, separated by strategy type (TA, FA, OS, NS) and predictive horizon (Long-term, Daily). Sentiments are calculated from bullish and bearish message counts within each subgroup. The explanatory and control variables are defined as in Table 3. All regressions include trading day fixed effects. Standard errors are clustered by trading day, with corresponding t-statistics in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Return Predictability at Longer Horizons

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
$\text{Sentiment}_{i,t}^{TA}$	-0.060***	-0.056***	-0.016
	[-3.65]	[-3.65]	[-1.09]
$\text{Sentiment}_{i,t}^{FA}$	0.052***	-0.002	-0.001
	[3.88]	[-0.12]	[-0.08]
$\text{Sentiment}_{i,t}^{OS}$	-0.091***	-0.046***	-0.036**
	[-5.30]	[-2.73]	[-2.28]
$\text{Sentiment}_{i,t}^{NS}$	-0.027**	-0.001	0.007
	[-2.37]	[-0.12]	[0.64]
$\text{Attention}_{i,t}$	-0.164***	-0.083***	-0.065***
	[-8.91]	[-6.05]	[-5.34]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

Panel B. Return Predictability by Strategy and Horizon

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
$\text{Sentiment}_{i,t}^{TA,Longterm}$	-0.045*	-0.049*	-0.016
	[-1.71]	[-1.95]	[-0.67]
$\text{Sentiment}_{i,t}^{FA,Longterm}$	0.051***	0.007	-0.030
	[2.66]	[0.37]	[-1.62]
$\text{Sentiment}_{i,t}^{OS,Longterm}$	-0.028	0.011	-0.066***
	[-1.00]	[0.38]	[-2.70]
$\text{Sentiment}_{i,t}^{TA,Daily}$	-0.149***	-0.075***	-0.033*
	[-6.50]	[-3.54]	[-1.65]
$\text{Sentiment}_{i,t}^{FA,Daily}$	-0.226***	-0.092***	-0.068**
	[-6.45]	[-2.93]	[-2.17]
$\text{Sentiment}_{i,t}^{OS,Daily}$	-0.237***	-0.086***	-0.062**
	[-8.53]	[-3.37]	[-2.50]
$\text{Attention}_{i,t}$	-0.153***	-0.078***	-0.058***
	[-8.54]	[-5.87]	[-4.96]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

Table 5: **Return Predictability by Investor Experience**

This table reports regression results examining whether investor experience affects the predictive power of retail investor sentiments on stock returns. Sentiment measures are categorized by trading strategy types (Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), Non-Strategy (NS)) and interacted with investor experience, measured as the fraction of messages posted by users self-identified as professional investors. All regressions include trading-day fixed effects. Standard errors are clustered by trading day, and t-statistics are reported in brackets. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
$\text{Sentiment}_{i,t}^{TA}$	-0.087***	-0.089***	-0.032
	[-4.00]	[-4.37]	[-1.60]
$\text{Sentiment}_{i,t}^{FA}$	0.047***	0.007	0.001
	[2.77]	[0.40]	[0.04]
$\text{Sentiment}_{i,t}^{OS}$	-0.116***	-0.063***	-0.055***
	[-5.13]	[-2.89]	[-2.64]
$\text{Sentiment}_{i,t}^{NS}$	-0.031**	-0.007	0.016
	[-2.31]	[-0.47]	[1.20]
$\text{Sentiment}_{i,t}^{TA} \times \text{Fraction of Messages by Professional}_{i,t}$	0.107***	0.128***	0.063*
	[3.03]	[3.78]	[1.95]
$\text{Sentiment}_{i,t}^{FA} \times \text{Fraction of Messages by Professional}_{i,t}$	0.027	-0.026	-0.004
	[0.75]	[-0.77]	[-0.11]
$\text{Sentiment}_{i,t}^{OS} \times \text{Fraction of Messages by Professional}_{i,t}$	0.118***	0.084**	0.085**
	[2.90]	[2.15]	[2.24]
$\text{Sentiment}_{i,t}^{NS} \times \text{Fraction of Messages by Professional}_{i,t}$	0.035	0.032	-0.039
	[1.26]	[1.15]	[-1.45]
$\text{Fraction of Messages by Professional}_{i,t}$	0.054**	0.017	0.027
	[2.27]	[0.71]	[1.21]
$\text{Attention}_{i,t}$	-0.161***	-0.081***	-0.064***
	[-8.78]	[-5.94]	[-5.25]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

Table 6: **Event Study: GameStop Short Squeeze**

This table reports results from an event study analyzing the impact of the GameStop (GME) short squeeze event on the informativeness of retail investor sentiments by different strategy types. The analysis covers a period of 180 trading days before and after January 12, 2021, the onset date of the GME short squeeze. Sentiment measures for each investment strategy category (Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), Non-Strategy (NS)) are interacted with an indicator for the post-GME period (January 14, 2021, onward). The explanatory and control variables are defined as in Table 3. All regressions include trading-day fixed effects. Standard errors are clustered by trading day, with corresponding t-statistics reported in brackets. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	$\text{Return}_{i,t+1\to t+5}$ (%)		$\text{Return}_{i,t+6\to t+10}$ (%)		$\text{Return}_{i,t+11\to t+15}$ (%)	
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Sentiment}_{i,t}^{TA}$	-0.060***	0.284***	-0.056***	0.252***	-0.016	0.144
	[-3.65]	[3.03]	[-3.65]	[2.64]	[-1.09]	[1.62]
$\text{Sentiment}_{i,t}^{FA}$	0.052***	0.132**	-0.002	0.024	-0.001	-0.061
	[3.88]	[2.15]	[-0.12]	[0.40]	[-0.08]	[-0.89]
$\text{Sentiment}_{i,t}^{OS}$	-0.091***	-0.019	-0.046***	0.238***	-0.036**	0.281***
	[-5.30]	[-0.23]	[-2.73]	[2.64]	[-2.28]	[3.37]
$\text{Sentiment}_{i,t}^{NS}$	-0.027**	-0.045	-0.001	0.138**	0.007	0.089
	[-2.37]	[-0.75]	[-0.12]	[2.04]	[0.64]	[1.39]
$\text{Sentiment}_{i,t}^{TA} \times$ Post-GameStop Episode		-0.455***		-0.468***		-0.295**
		[-3.20]		[-3.39]		[-2.12]
$\text{Sentiment}_{i,t}^{FA} \times$ Post-GameStop Episode		-0.176*		-0.173*		-0.043
		[-1.83]		[-1.84]		[-0.44]
$\text{Sentiment}_{i,t}^{OS} \times$ Post-GameStop Episode		-0.219*		-0.455***		-0.491***
		[-1.67]		[-3.28]		[-3.85]
$\text{Sentiment}_{i,t}^{NS} \times$ Post-GameStop Episode		-0.113		-0.345***		-0.220**
		[-1.25]		[-3.65]		[-2.30]
$\text{Attention}_{i,t}$	-0.164***	-0.026	-0.083***	0.046	-0.065***	0.119
	[-8.91]	[-0.13]	[-6.05]	[0.32]	[-5.34]	[1.00]
Stock Characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	2,972,286	530,098	2,970,213	529,747	2,968,174	529,368
R^2	0.105	0.082	0.108	0.091	0.112	0.087

Table 7: **Retail Investor Sentiment and Market Order Imbalances**

Panel A examines the contemporaneous relation between retail investor sentiment revealed by Stocktwits messages (classified by investment strategy type: Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), and Non-Strategy (NS)) and aggregate retail market order imbalances (*OIB*). We separately measure sentiment during intraday trading hours (IN) and overnight periods (OV), estimating the following panel regression at the stock-day level:

$$OIB_{i,t} = \beta_1 Sentiment_{i,t}^{type,IN} + \beta_2 Sentiment_{i,t}^{type,OV} + \beta_3 Attention_{i,t} + \gamma \boldsymbol{X_{i,t}} + \delta_t + \epsilon_{i,t}.$$

The dependent variable *OIB* is measured using two methods, following Boehmer et al. (2021) (*BJZZ*) and Barber et al. (2022) (*BJHOS*). Panel B investigates the informativeness of sentiment-driven *OIB* for next-day stock returns. Specifically, we first decompose *OIB* into strategy-specific components by regressing *OIB* on intraday sentiments of all four strategy types simultaneously. These sentiment-driven *OIB* components are then used to predict next-day returns. The explanatory and control variables are defined as in Table 3. All specifications include trading-day fixed effects. Standard errors are clustered by stock and trading day in Panel A and by stock in Panel B. Corresponding t-statistics are reported in brackets. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Contemptuous Relation Between Stocktwits Sentiments and Retail *OIB*

	$OIB_{i,t}^{BJZZ}$ (%)			$OIB_{i,t}^{BHJOS}$ (%)		
	(1)	(2)	(3)	(4)	(5)	(6)
$Sentiment_{i,t}^{TA,IN}$	0.728***		0.732***	0.975***		0.944***
	[17.48]		[17.64]	[24.97]		[24.52]
$Sentiment_{i,t}^{FA,IN}$	0.566***		0.571***	0.612***		0.586***
	[15.02]		[15.31]	[16.33]		[15.96]
$Sentiment_{i,t}^{OS,IN}$	0.719***		0.723***	0.792***		0.758***
	[19.06]		[19.35]	[22.89]		[22.67]
$Sentiment_{i,t}^{NS,IN}$	0.485***		0.493***	0.634***		0.619***
	[13.57]		[13.93]	[18.20]		[18.00]
$Sentiment_{i,t}^{TA,OV}$		0.172***	0.002		0.447***	0.249***
		[4.58]	[0.05]		[12.60]	[7.32]
$Sentiment_{i,t}^{FA,OV}$		0.163***	0.004		0.312***	0.131***
		[4.30]	[0.10]		[8.62]	[3.70]
$Sentiment_{i,t}^{OS,OV}$		0.283***	0.067*		0.340***	0.092**
		[6.84]	[1.69]		[9.03]	[2.55]
$Sentiment_{i,t}^{NS,OV}$		0.028	-0.131***		0.187***	0.002
		[0.78]	[-3.65]		[5.39]	[0.05]
$Attention_{i,t}$	0.385***	0.410**	0.384***	0.481***	0.503***	0.474***
	[2.75]	[2.42]	[2.74]	[3.38]	[2.97]	[3.47]
Stock Characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	2,974,934	2,974,934	2,974,934	2,974,934	2,974,934	2,974,934
R^2	0.009	0.008	0.009	0.012	0.011	0.012

Panel B. Informativeness of Strategy-Specific Retail *OIB*

	Return$_{i,t+1}$ (%)	
	(1)	(2)
$OIB_{i,t}^{BJZZ,TA}$	-3.557***	
	[-3.17]	
$OIB_{i,t}^{BJZZ,FA}$	5.555***	
	[2.98]	
$OIB_{i,t}^{BJZZ,OS}$	-6.884***	
	[-4.58]	
$OIB_{i,t}^{BJZZ,NS}$	-0.042	
	[-0.02]	
$OIB_{i,t}^{BJZZ,Resid}$	0.173***	
	[13.62]	
$OIB_{i,t}^{BHJOS,TA}$		-2.524***
		[-3.15]
$OIB_{i,t}^{BHJOS,FA}$		4.255***
		[2.99]
$OIB_{i,t}^{BHJOS,OS}$		-5.428***
		[-4.57]
$OIB_{i,t}^{BHJOS,NS}$		-0.038
		[-0.03]
$OIB_{i,t}^{BHJOS,Resid}$		0.243***
		[18.25]
Attention$_{i,t}$	-0.056***	-0.056***
	[-5.50]	[-5.55]
Stock Characteristics	Yes	Yes
Lagged Returns	Yes	Yes
Time FEs	Yes	Yes
N	2,974,304	2,974,304
R^2	0.089	0.089

Table 8: **Robinhood Herding and StockTwits Sentiments**

This table investigates the contemporaneous relation between retail investor herding on Robinhood and sentiments derived from StockTwits messages, categorized by strategy type: Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), and Non-Strategy (NS). We estimate the following panel regression at the stock-day level:

$$RH\ Herding_{i,t} = \beta_1 Sentiment_{i,t}^{type} + \beta_2 Attention_{i,t} + \gamma X_{i,t} + \delta_t + \epsilon_{i,t},$$

where the dependent variable, $RH\ Herding_{i,t}$, is an indicator equal to one if stock i is among the top 10 stocks ranked by daily percentage growth in Robinhood users holding the stock, provided at least 100 users held the stock at the end of day $t-1$, as defined in Barber et al. (2022). Robinhood user account data span May 2018 through August 2020, sourced from RobinTrack. The explanatory and control variables are defined as in Table 3. All regressions include trading-day fixed effects. Standard errors are clustered by stock and trading day, and t-statistics are presented in brackets. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	RH Herding$_{i,t}$ (%)				
	(1)	(2)	(3)	(4)	(5)
Sentiment$_{i,t}^{TA}$	0.137***				0.117***
	[6.40]				[6.18]
Sentiment$_{i,t}^{FA}$		0.090***			0.066***
		[4.88]			[4.09]
Sentiment$_{i,t}^{OS}$			0.090***		0.066***
			[3.86]		[3.13]
Sentiment$_{i,t}^{NS}$				0.035***	0.019
				[2.76]	[1.62]
Attention$_{i,t}$	1.183***	1.185***	1.185***	1.188***	1.180***
	[3.70]	[3.70]	[3.70]	[3.70]	[3.71]
Stock Characteristics	Yes	Yes	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes	Yes	Yes
N	554,877	554,877	554,877	554,877	554,877
R^2	0.013	0.013	0.013	0.013	0.013

7. Appendix A

Table A.1: **Examples of GPT Responses**

No.	Message	Ticker	Score	Indicators
1	$IOVA Biotechnology Company, Phase 2, Hammer, Support Line, Oversold, JMP Securities $38, Q4: Institutional Bought $77M, Sold $13M, Speculation Trade, Entry: Above $24	IOVA	2	Hammer, Support Line, Oversold
2	$CVS if it can hold firmly above $106 will signal entry at the close as well. Stops tight at $104	CVS	2	Support Level, Stop Loss
3	$RETA 10 wk SMA has caught up. $300 stock btw, Livermore's finest	RETA	2	10 wk SMA
4	RT @mentholatum $AAPL the oversold compression on AAPL will release... another $50 up day maybe.... when????.... Someday soon// Bold call	AAPL	2	Oversold Compression
5	$AAPL next retracement $100.36 which is 38.2% of the move down. should be coming within next hr	AAPL	2	Fibonacci Retracement
6	$ACOR Acorda Therapeutics (ACOR, $8.65) was this week's top stock market loser, declining -10%. Expect a Downtrend reversal	ACOR	1	Downtrend Reversal
7	$META Bout to break the big $100 level then breakdown further.	META	1	Breakdown
8	$SAIC Science Applications International Corporation (SAIC) has been systematically hitting all-time highs in the last 10 days. Science Applications International Corporation (SAIC) price climbed on Wednesday a 2.17% ending at $103.10 and marking the n	SAIC	1	All-time Highs
9	$PETS new retail shorts probably got in at 35 or lower, this will fly on short covering above $38.50ish when most down over 10%	PETS	1	Short Covering
10	10:27:29 AM Makes fresh HOD $CARA $19.55 +12.2% ON 1,400K VOL (ISW Pre-Market Watch/Scan)	CARA	1	HOD, Volume
11	$TSLA added more under $890 ... well it has been while since last time I played with TSLA... I just love how their earning growing and what ELON said... I still expect volatile days but worth to start adding... GL	TSLA	0	-
12	$MSFT Lmaooo you bears are dumb as shit. I sold all my Bitcoin to buy shares at $275 hand over fist.	MSFT	0	-
13	$MU I picked up some of the $25s for a punt...Company is undervalued massively...if they deliver, this soars > 15%.	MU	0	-
14	$ETSY at $13.66 - Sell Stock Market Alert sent at 10:14 AM ET #stocks	ETSY	0	-

Table A.2: **Variable Definitions**

Variable	Definition
$\text{Usage}_{j,i,t,n}^{type}$	Fraction of total messages posted by investor j on stock i in day t that are classified as being related to investment strategy type.
$\text{Sentiment}_{i,t}$	The difference in the number of bullish and bearish messages to the sum of bullish and bearish messages on stock i posted on Stocktwits in day t, $\frac{N_{i,t}^{Bullish}-N_{i,t}^{Bearish}}{N_{i,t}^{Bullish}+N_{i,t}^{Bearish}}$ following Cookson and Niessner (2020).
$\text{Sentiment}_{i,t}^{TA}$	Sentiment calculated using messages related to technical analysis.
$\text{Sentiment}_{i,t}^{FA}$	Sentiment calculated using messages related to fundamental analysis.
$\text{Sentiment}_{i,t}^{OS}$	Sentiment calculated using messages that are not related to other strategies.
$\text{Sentiment}_{i,t}^{NS}$	Sentiment calculated using messages that are not related to technical analysis, fundamental analysis, or other strategies.
$\text{Technical Investor}_{j}$	Dummy variable equal to one if investor i's self-reported approach is "Technical" or "Momentum".
$\text{Swing or Day Trader}_{j}$	Dummy variable equal to one if investor i's self-reported holding period is "Swing Trader" or "Day Trader".
$\text{Long-Term Investor}_{j}$	Dummy variable equal to one if investor i's self-reported holding period is "Long-Term Investor".
Professional_{j}	Dummy variable equal to one if investor i's self-reported experience is "Professional".
$\text{Attention}_{i,t}$	A measure of StockTwits users' attention on stock i in week t, defined as the number of messages on stock i divided by the total number of messages across all stocks, i.e., $Attention_{i,t} = \frac{\#Messages_{i,t}}{\sum_i \#Messages_{i,t}}$ (Cookson et al., 2024b).
$\text{OIB}_{i,t}$	Retail marketable volume imbalance on stock i in week t following Boehmer et al. (2021) (BJZZ) or Barber et al. (2023) (BHJOS).
$\text{RH Herding}_{i,t}$	Indicator for top 1% of positive Robinhood user change ratio in week t and a minimum of 100 users at the end of week $t-1$ following Barber et al. (2022).
$\text{MAX}_{i,t}$	Maximum one-day return in the prior month.
$\text{Abnormal Turnover}_{i,t}$	Measure of abnormal trading volume, $\log(1+Turnover_{i,t}) - \log(1+\frac{1}{4}\sum_{h=1}^{4} Turnover_{i,t-h})$.
$\text{Abnormal News}_{i,t}$	Measure of abnormal volume of news articles reported in Ravenpack, $\log(\#News_{i,t}) - \log(1+\frac{1}{4}\sum_{h=1}^{4} \#News_{i,t-h})$.
$\text{Market Capitalization}_{i,t}$	Market capitalization.
$\text{Book-to-Market}_{i,t}$	Ratio of book value to market value.
$\text{Asset Growth}_{i,t}$	Growth rate of annual total assets.
$\text{Gross Profits-to-Asset}_{i,t}$	Ratio of gross profits to total assets.
$\text{Number of Analysts}_{i,t}$	Number of IBES equity analysts covering stock i.
$\text{Institutional Ownership}_{i,t}$	Fraction of shares outstanding held by 13F institutional investors.

Table A.3: **Predicting Returns using Fama-MacBeth Regressions**

This table presents the return predictability results with Stocktwits sentiments using Fama-MacBeth regressions.

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
$\text{Sentiment}_{i,t}^{TA}$	-0.047***	-0.055***	-0.016
	[-3.51]	[-4.35]	[-1.31]
$\text{Sentiment}_{i,t}^{FA}$	0.042***	-0.012	-0.011
	[3.36]	[-1.00]	[-0.93]
$\text{Sentiment}_{i,t}^{OS}$	-0.064***	-0.032**	-0.034***
	[-4.41]	[-2.27]	[-2.61]
$\text{Sentiment}_{i,t}^{NS}$	-0.026**	-0.003	0.004
	[-2.54]	[-0.26]	[0.37]
$\text{Attention}_{i,t}$	-0.250***	-0.120***	-0.102***
	[-7.71]	[-4.53]	[-4.33]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.107	0.099	0.099

Table A.4: **Predicting DGTW Excess Returns with Retail Sentiments**

This table examines the predictability of retail investor sentiments on DGTW excess returns.

	$DGTW_{i,t+1\to t+5}$ (%)	$DGTW_{i,t+6\to t+10}$ (%)	$DGTW_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
$Sentiment^{TA}_{i,t}$	-0.059***	-0.059***	-0.024*
	[-4.17]	[-4.37]	[-1.81]
$Sentiment^{FA}_{i,t}$	0.048***	0.000	-0.002
	[3.97]	[0.04]	[-0.18]
$Sentiment^{OS}_{i,t}$	-0.063***	-0.031**	-0.038***
	[-4.10]	[-2.16]	[-2.70]
$Sentiment^{NS}_{i,t}$	-0.025**	0.002	0.005
	[-2.43]	[0.19]	[0.54]
$Attention_{i,t}$	-0.097***	-0.047***	-0.023**
	[-6.14]	[-4.19]	[-2.45]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes
N	2,662,982	2,660,293	2,657,567
R^2	0.009	0.009	0.009

Table A.5: **Return Predictability on Stock-Day with at least 10 messages**

This table revisits the analysis of daily return predictability using the sample where we restrict to stock-day with at least 10 Stocktwits messages as in Cookson et al. (2024b).

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
Sentiment^{TA}	-0.077**	-0.085**	0.047
	[-2.44]	[-2.58]	[1.62]
Sentiment^{FA}	0.194***	0.071*	0.039
	[4.99]	[1.93]	[1.08]
Sentiment^{OS}	-0.172***	-0.028	-0.038
	[-5.03]	[-0.92]	[-1.36]
Sentiment^{NS}	-0.000	0.013	-0.008
	[-0.00]	[0.26]	[-0.19]
Attention	-0.081***	-0.046***	-0.038***
	[-6.56]	[-4.85]	[-4.48]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes
N	638,965	638,570	638,198
R^2	0.107	0.114	0.128

Table A.6: **Predicting Returns with Cookson and Niessner (2020) Sentiment Measures**

This table revisits the daily return predictability results in Table 3 using the alternative sentiment measures computed based on the maximum entropy method proposed in Cookson and Niessner (2020).

	$\text{Return}_{i,t+}$ (%)				
	(1)	(2)	(3)	(4)	(5)
$\text{Sentiment}^{TA}_{i,t}$	-0.024***				-0.022***
	[-3.13]				[-3.18]
$\text{Sentiment}^{FA}_{i,t}$		0.010			0.016***
		[1.60]			[2.69]
$\text{Sentiment}^{OS}_{i,t}$			-0.032***		-0.031***
			[-3.91]		[-4.03]
$\text{Sentiment}^{NS}_{i,t}$				-0.006	-0.002
				[-0.99]	[-0.40]
$\text{Attention}_{i,t}$	-0.055***	-0.056***	-0.055***	-0.056***	-0.055***
	[-5.44]	[-5.50]	[-5.42]	[-5.48]	[-5.41]
Stock Characteristics	Yes	Yes	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes	Yes	Yes
N	2,974,304	2,974,304	2,974,304	2,974,304	2,974,304
R^2	0.089	0.089	0.089	0.089	0.089

Table A.7: **Predicting Returns at Longer Horizon with Sentiment-Driven Retail Market Order Flows**

This table examines the longer horizon return predictability of retail market order imbalance (*OIB*) driven by strategy-specific sentiments as in Table 7.

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
$\text{OIB}_{i,t}^{BJZZ,TA}$	-13.473***	-8.357***	-3.946
	[-4.88]	[-3.25]	[-1.63]
$\text{OIB}_{i,t}^{BJZZ,FA}$	10.292**	-4.616	-6.052
	[2.40]	[-1.13]	[-1.45]
$\text{OIB}_{i,t}^{BJZZ,OS}$	-20.966***	-9.416***	-5.834*
	[-5.77]	[-2.77]	[-1.74]
$\text{OIB}_{i,t}^{BJZZ,NS}$	-17.199***	-4.204	1.325
	[-3.42]	[-0.83]	[0.27]
$\text{OIB}_{i,t}^{BJZZ,Resid}$	0.261***	0.085***	0.030
	[9.42]	[3.02]	[1.06]
$\text{Attention}_{i,t}$	-0.162***	-0.082***	-0.064***
	[-8.89]	[-6.01]	[-5.31]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
$\text{OIB}_{i,t}^{BHJOS,TA}$	-9.579***	-5.949***	-2.808
	[-4.87]	[-3.25]	[-1.63]
$\text{OIB}_{i,t}^{BHJOS,FA}$	7.868**	-3.513	-4.622
	[2.40]	[-1.13]	[-1.45]
$\text{OIB}_{i,t}^{BHJOS,OS}$	-16.554***	-7.428***	-4.608*
	[-5.77]	[-2.77]	[-1.74]
$\text{OIB}_{i,t}^{BHJOS,NS}$	-9.962***	-2.415	0.758
	[-3.43]	[-0.83]	[0.27]
$\text{OIB}_{i,t}^{BHJOS,Resid}$	0.426***	0.081***	0.051
	[14.27]	[2.70]	[1.58]
$\text{Attention}_{i,t}$	-0.163***	-0.082***	-0.064***
	[-8.94]	[-6.02]	[-5.32]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

Table A.8: **252-Day Rolling Estimated Sentiment-Driven Retail Market Order Flows**

This table investigates the return predictability of sentiment-driven retail market order imbalances, which are estimated using an alternative 252-day rolling-window method and decomposed into strategy-specific components by regressing *OIB* on intraday sentiments of all four strategy types. These sentiment-driven *OIB* components, which are not estimated only using the past information, are then used to predict next-day returns.

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
$\text{OIB}_{i,t}^{BJZZ,TA}$	-15.594***	-8.691***	-5.174**
	[-6.15]	[-3.63]	[-2.29]
$\text{OIB}_{i,t}^{BJZZ,FA}$	4.661*	-5.195**	-4.735*
	[1.69]	[-2.00]	[-1.82]
$\text{OIB}_{i,t}^{BJZZ,OS}$	-15.849***	-7.543***	-5.367**
	[-6.54]	[-3.36]	[-2.39]
$\text{OIB}_{i,t}^{BJZZ,NS}$	-3.385	1.023	1.736
	[-1.49]	[0.46]	[0.78]
$\text{OIB}_{i,t}^{BJZZ,Resid}$	0.260***	0.085***	0.030
	[9.41]	[3.01]	[1.07]
$\text{Attention}_{i,t}$	-0.163***	-0.082***	-0.064***
	[-8.90]	[-6.01]	[-5.29]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
$\text{OIB}_{i,t}^{BHJOS,TA}$	-8.414***	-5.232***	-1.820
	[-5.54]	[-3.62]	[-1.33]
$\text{OIB}_{i,t}^{BHJOS,FA}$	6.512***	-2.299	-2.071
	[2.66]	[-0.99]	[-0.89]
$\text{OIB}_{i,t}^{BHJOS,OS}$	-14.375***	-7.796***	-5.787***
	[-7.09]	[-4.22]	[-3.10]
$\text{OIB}_{i,t}^{BHJOS,NS}$	-2.834*	-0.363	0.900
	[-1.71]	[-0.22]	[0.56]
$\text{OIB}_{i,t}^{BHJOS,Resid}$	0.425***	0.080***	0.050
	[14.22]	[2.69]	[1.57]
$\text{Attention}_{i,t}$	-0.164***	-0.083***	-0.065***
	[-8.97]	[-6.03]	[-5.33]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
Time FEs	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112