Wisdom or Whims? Decoding Retail Strategies with Social Media and AI*

Shuaiyu Chen† Lin Peng ‡ Dexin Zhou§

August 28, 2025

Abstract

We use rich social media data and large language models to classify retail investors' trading strategies and show that strategy adoption is highly dynamic, responding to news, past performance, and social feedback. Sentiment from Fundamental Analysis posts predicts positive future returns, whereas sentiments from Technical Analysis and other speculative strategies posts predict negative returns. Technical sentiment is especially correlated with retail buying, most prominently among Robinhood investors. Crucially, retail order flows are informative only when linked to fundamental-analysis sentiment. These results illuminate how retail investors form strategies and the conditions under which they enhance or distort market outcomes.

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

*First draft: September 26, 2024. The paper was previously circulated under the title "Wisdom or Whims? Decoding Investor Trading Strategies with Large Language Models." We appreciate suggestions from Tony Cookson (discussant), Zhi Da, Chukwuma Dim, Winston Dou, David Hirshleifer, Byoung-Hyun Hwang, Celine Fei (discussant), Clifton Green, Marina Niessner, Ron Kaniel, John Shim, Dacheng Xiu, Zonghao Yang (discussant), Chishen Wei (discussant), and seminar participants at Baruch College, Purdue University, Binghamton University, George Washington University, McMaster University, Rochester Institute of Technology, University of Central Florida, University of Notre Dame, Tsinghua University (PBCSF), 4th Hong Kong Conference for Fintech, AI, and Big Data in Business, ABFER, CICF, Gen AI Conference, Oxford-Man Machine Learning in Finance Conference, TUM Workshop on NLP and GenAI, USC Conference on Social and Behavioral Finance, and Wolfe Research AI in Finance Conference.

†chens@darden.uva.edu, Darden School of Business, University of Virginia.

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

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

1. Introduction

The recent rise of fintech brokerage platforms and social media sites has stimulated significant increases in stock market participation by retail investors.[1] This dramatic expansion raises questions about the implications of retail investors' growing presence. Prior studies have shown that retail investors can act either as noise traders susceptible to behavioral biases or as informed traders whose orders predict returns and provide liquidity.[2] Yet, retail investors' decision-making processes that lead to these diverging outcomes remain insufficiently understood.

To address this gap, we leverage rich, real-time social media data and large language models (LLMs) to analyze how retail investors select strategies, adjust them over time, and what these choices imply for trading outcomes. This approach is crucial, as it enables us to accurately extract complex financial concepts from social media messages, which tend to be highly colloquial and fragmented.

A key insight from our analysis is that investors have considerable flexibility in strategy use, dynamically shifting across approaches as conditions evolve and new information arrives. Figure 1 illustrates this behavior with screenshots from StockTwits. The figure features a self-declared technical strategist, *ACInvestorBlog*, who in one post discusses earnings potential and valuations, and in another relies on technical indicators such as support and resistance levels and volume patterns. These posts highlight that individual investors often shift between fundamental and technical approaches rather than adhering rigidly to a single approach.

Building on this observation, we conduct a systematic analysis of when investors adopt a particular approach, how and when they switch, and how these choices map into trading and price outcomes. We apply LLMs to a dataset of 96 million

[1] By 2021, retail trading in the U.S. accounts for almost as much volume as mutual funds and hedge funds combined. Furthermore, at the start of 2023, U.S. retail investor daily inflows reached a record-breaking high of $1.5 billion—more than double the pre-2019 figure of just over $600 million. Source: K. Martin and R. Wigglesworth, *Rise of the retail army: the amateur traders transforming markets*, Financial Times, March 9 2021; and P. Rao, *Charted: U.S. Retail Investor Inflows (2014–2023)*, Visual Capitalist, November 5, 2023.

[2] See, e.g., Barber and Odean (2000); Kumar and Lee (2006); Barber and Odean (2008); Barber et al. (2022); Bryzgalova et al. (2023); Barber et al. (2023b) for the former view, and Kaniel et al. (2008); Kelley and Tetlock (2013); Boehmer et al. (2021); Welch (2022) for the latter.

StockTwits messages covering roughly 7,800 stocks and 840,000 users from January 2010 to June 2023. StockTwits, one of the largest investment social media platforms, has become an increasingly important data source in recent studies examining retail investor attention, beliefs, and sentiments (e.g., Cookson and Niessner, 2020; Cookson et al., 2023, 2024b).

Using LLMs, we classify each message into four categories based on the trading strategies reflected: technical analysis (TA), fundamental analysis (FA), other strategies (OS, such as options trading), and no strategy (NS). We find that 31% of all messages explicitly reference an investment strategy. Among these, 44% reference FA, 28% reference TA, and 28% reference OS.[3]

Our classification highlights that strategy use is highly malleable, with much of the variation unexplained by observable investor and stock characteristics. While correlated with known investor attributes—for example, TA is more common among self-declared momentum or short-horizon traders, whereas FA is more common among long-term investors—these factors explain only a limited share of the variation. Self-declared attributes account for a single-digit percentage of message-level variation, and even with time and investor-stock fixed effects, the R-squared reaches only about 30%. The substantial within-investor-stock variation that remains unexplained suggests that strategy use is a dynamic, context-dependent choice rather than a fixed trait.

Motivated by this, we examine the factors influencing dynamic strategy adoption and uncover three novel patterns. First, information availability matters: following firm news, investors rely more on FA and less on TA and OS. In particular, they are substantially more likely to discuss FA, by 28% of the mean, following news releases. This is consistent with the view that investors are attentive to public news and increase their use of FA when fundamental information are available.

Second, performance feedback shapes behavior: poor returns from a strategy in-

[3]Messages may be classified into more than one category, but fewer than 1% mention multiple strategies. The remaining 69% include memes, catchphrases like "to the moon," or other content unrelated to specific strategies.

duce switching, while strong returns from alternative strategies attract investors away. This pattern is consistent with the behavior modeled in Barberis et al. (2018), where investors weigh both price and valuation signals and shift their emphasis between the signals over time. Our evidence suggests that such switching is endogenous and responsive to investors' own performance, pointing to opportunities for future models to incorporate this dynamic explicitly.

Third, social interactions matter: positive feedback (likes) on posts employing a strategy increases its continued use, while positive feedback on posts using alternative strategies reduces reliance on the focal strategy. These findings support the view in Hirshleifer (2020) that social transmission biases favor folk models that are more heavily cued in the environment. Moreover, in many cases, investors' strategy choices can be more heavily influenced by social feedback than their own past performance, underscoring the central role of social dynamics in shaping investor behavior.[4]

Having established these dynamics, we next examine how strategy use translates into market consequences, focusing on return predictability, retail order flows, informativeness, and speculative behavior.

We first study performance implications. Stocks with bullish TA or OS sentiment earn significantly lower next-day returns, whereas stocks with bullish FA sentiment outperform. This suggests that, on average, technical and speculative strategies perform poorly in the short run, while fundamental analysis is informative about future returns. The economic magnitudes are sizable: contrarian strategies that bet against TA or OS sentiment yield annualized returns of about 10%, while following FA sentiment yields about 7.5%.[5] Furthermore, investor sophistication, measured by the fraction of messages posted by users self-reporting

[4]For example, users' TA strategy choices are 85% more responsive to a low number of "likes" in past TA messages than to poor resulted from their past TA usage.

[5]To achieve these returns, one must trade contrary to TA and OS sentiment, which are negatively related to future returns. These estimates do not incorporate transaction costs, and the high turnover implied by such strategies would likely erode realizable returns. Our objective is therefore not to propose implementable trading strategies but to shed light on retail investor behavior on social media and the market frictions it creates.

professional experience, modulates these relationships. Specifically, the underperformance of TA and OS sentiment is less pronounced when professional participation is higher, whereas the predictive power of FA remains largely unaffected.

Next, we link strategy-specific sentiment on social media to retail order flows. We first test whether sentiment expressed on StockTwits is reflected in retail investors' trading activity by regressing retail market order imbalance measures on strategy-specific sentiment. We find that sentiments across all strategies are positively associated with net retail buying, indicating that StockTwits provides a real-time window into retail trading activity.

We then turn to the question of informativeness, one of the central issues in the literature discussed earlier: does retail trading represent "smart money" that conveys valuable information, or "noise" driven by behavioral biases? Our approach provides new insights into this debate by allowing for heterogeneity and flexibility in retail investors' strategies. Decomposing retail order flow imbalance into components attributable to each social media-based strategy, we find that TA- and OS-related flows weaken the ability of retail order imbalance to predict future returns, while FA-related flow strengthens it.

These results add nuance to the evidence in Boehmer et al. (2021) and Barber et al. (2023a), who document the informativeness of retail order flows. We go further by showing that a sizable share of this information might have been publicly articulated in social media rather than concealed in trading activity. This willingness of retail investors to disclose information online challenges the traditional view that informed traders seek secrecy and opens new directions for theories of price discovery and the value of private information.

Finally, we focus on Robinhood investors, who played a central role in speculative episodes such as the "meme stock mania." As shown in Barber et al. (2022), these investors engage in highly attention-driven trading and exhibit intense buying episodes ("herding") that are typically followed by sharp negative returns. Linking strategy use to these episodes, we find that bullish TA sentiment is most strongly

associated with Robinhood herding. This evidence suggests that crowding on technical signals, which are widely promoted on social media platforms, can be a key driver of speculative surges. The mechanism resonates with Stein (2009): when traders cannot observe how many others follow the same signals, coordination failures push prices away from fundamentals. Among retail investors, the effect is magnified by the salience of technical signals and their limited ability to short, resulting in concentrated buying pressure and subsequent reversals.

Taken together, these findings provide new insights into retail investors' behavior and its asset pricing implications. Retail order flows informed by fundamental analysis incorporate valuable information and enhance market efficiency, whereas crowding on technical signals amplifies inefficiencies and fuels reversals.

Our novel two-step procedure, which first uses a state-of-the-art LLM to generate labeled examples and then fine-tunes a smaller model for large-scale classification, substantially reduces computational costs while preserving accuracy. Applied to rich social media data, this methodology provides a broader framework for studying retail investors and their role in financial markets, including how they interact with more sophisticated participants and how such interactions affect efficiency and stability. For example, Chen et al. (2025) show that the primary profits of AI-powered strategies come from trading against retail TA sentiment, a result closely aligned with the theoretical settings of Dou et al. (2024). Looking ahead, this framework opens new avenues for understanding when retail participation enhances price discovery and when it destabilizes markets, a distinction with direct implications for platform regulation, investor protection, and the broader debate on how social media and AI shape market stability.

Our study contributes to several strands of literature. First, it extends the growing research on retail trading and its informativeness. The rise of zero-cost trading platforms has brought a large influx of new investors into financial markets, reigniting interest in the role of retail traders (e.g., Barber et al., 2022; Welch, 2022; Eaton et al., 2022). Yet the investment approaches these investors employ remain insuf-

ficiently understood.[6] Our contribution is to advance this literature by extracting investor strategies directly from their own words. We are able to characterize how investors' beliefs and strategy choices evolve dynamically in response to information arrival, performance cues, and social feedback. Our analysis further shows that the informativeness of retail sentiment and order flow differs sharply across strategies discussed on social media.

Second, our study contributes to the literature on how investors employ fundamental and technical analysis.[7] Prior work, relying on static, self-declared investment approaches, shows that heterogeneity in these strategies explains a significant share of investor disagreement (Cookson and Niessner, 2020). We advance this research by showing that strategy use is not a fixed user attribute but instead evolves dynamically. This finding of dynamic switching provides an empirical microfoundation for behavioral models, such as Barberis et al. (2018), that exogenously assume investors waver between different signals over time. Our evidence of an endogenous process thus opens the door for future research to explore the implications of such dynamics for investor behavior and market outcomes.

Third, we contribute to the literature on social media and financial markets. Prior studies establish that social networks shape investor beliefs and affect stock returns and trading volume.[8] We build on this line of research by being the first to apply large language models (LLMs) to systematically extract trading strategies from social media data.[9] Our approach allows us to contribute to this liter-

[6]Characterizing retail investors' beliefs and investment strategies is challenging. Recent literature uses surveys to gain insights into investor beliefs. (e.g., 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). However, surveys are costly to conduct and cannot capture investors' thought processes in real time, limiting their ability to reveal how investors form beliefs and make trading decisions.

[7]For studies on fundamental analysis, see, for example, Porta et al. (1997) and Abarbanell and Bushee (1997). For technical analysis, see, e.g., 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), and more recent evidence in (e.g., Han et al., 2013, 2016; Jiang et al., 2023; Murray et al., 2024).

[8]See, for example, Antweiler and Frank (2004); Chen et al. (2014); Giannini et al. (2018, 2019); Cookson and Niessner (2020); Hirshleifer (2020); Han et al. (2022). See Cookson et al. (2024b) for a comprehensive review.

[9]Other recent applications of LLMs in economics and finance include Korinek (2023), Jiang et al. (2024a), Lopez-Lira and Tang (2023), and Huang et al. (2024).

ature in two ways. First, complementing prior evidence of confirmation bias and "echo chamber" effect in signal reception (e.g., Cookson et al., 2023), we show that social validation is another important determinant of investors' strategy choices, consistent with the social transmission bias model of Hirshleifer (2020). Second, our approach further uncovers new nuances in the relation between social media sentiment and returns, and in how social media sentiment moderates the informativeness of retail order flows.

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.[10] At the message level, our data include timestamps, textual content, and user-provided

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

sentiment labels ("bullish" or "bearish") when available. Additionally, we obtain user-level biographical characteristics self-reported by StockTwits users, including investment approach (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).

After merging with the CRSP stock universe, we apply several filters by following Cookson and Niessner (2020) and Cookson et al. (2024a) to ensure message validity and to focus 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. Our final message sample comprises 96,095,345 messages from 840,846 unique users for 7,834 stocks from January 2010 to June 2023.

2.2. Extracting Information from StockTwits Texts

In this section, we extract two key textual variables from StockTwits messages. We first show how large language models are used to identify the trading strategies referenced in individual posts. We then explain how sentiment is assigned to those messages without user-provided sentiment labels and further aggregated into stock-level, strategy-specific sentiment measures.

2.2.1. Identifying Trading Strategies from Messages

Identifying investment strategies from social media posts is a highly challenging task. StockTwits messages, like much social media content, are short, filled with abbreviations and colloquial expressions, and often contain non-standard spellings or variations. These features make it difficult to identify trading strategies using

dictionary-based approaches alone.

We address this challenge with a two-step classification approach that enables us to capture the likelihood that a message reflects a given strategy even when the language is noisy or indirect. In the first step, we apply the GPT-4 Turbo model to a small random sample of messages and generate a probability score for whether each message references a given trading strategy. In the second step, we use these outputs to fine-tune a smaller, more efficient BERT model on the full dataset. This approach, first proposed by Hinton (2015), is widely known as knowledge distillation (KD) in the machine learning literature.[11]

This methodology offers several advantages. GPT-4 Turbo classifications align closely with human judgment—for example, in a benchmark of 500 messages manually coded by a PhD research assistant, GPT-4 Turbo achieves a 90% correlation with the human labels. At the same time, applying GPT-4 Turbo to the entire dataset is highly costly, whereas BERT provides a much more efficient solution once fine-tuned.[12] Moreover, the two-step procedure outperforms natural alternatives such as a pure dictionary-based approach or an alternative two-step procedure that first uses user-declared investment approaches to train BERT. Appendix Section A.1 provides further details of these comparisons.

For each message, the two-step procedure evaluates whether it refers to one of four investment strategy categories: technical analysis, fundamental analysis, other strategies, or no strategy. Specifically, the model outputs a probability score for each strategy, and we classify the message as strategy-related if the score exceeds 95%.

[11] See Gu et al. (2023) for a recent review of KD and its applications in LLMs. We use the *bert-base-uncased* model, which has 110 million parameters. While large, this is modest compared to GPT-4 Turbo, rumored to have 1.7 trillion parameters. BERT has established itself as an efficient and capable tool for natural language processing tasks, including classification (Devlin et al., 2018). González-Carvajal and Garrido-Merchán (2020) show that BERT consistently outperforms traditional NLP tools that do not rely on deep learning.

[12] Our back-of-the-envelope calculation indicates that the total cost could exceed $500,000 to classify all messages.

Technical Analysis Strategy Classification Next, we illustrate our classification procedure by identifying messages that refer to technical analysis. We begin by randomly sampling $20{,}000$ messages from the full dataset, drawing $10{,}000$ from users with a self-declared technical investment approach and $10{,}000$ from other self-declared approaches. We then prompt GPT-4 Turbo to determine whether each message reflects 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": }.

Appendix Table A.1 provides examples of responses.[13] We then use GPT's outputs to fine-tune a BERT model (henceforth TA-BERT) to predict whether a message uses technical analysis. Cross-validation indicates that TechBERT achieves an F1 score of 0.83, which indicates a high level of performance. Because BERT has a drastically smaller parameter count than GPT-4 Turbo, TA-BERT can be run locally to generate probabilistic predictions of TA usage.

We next visualize TA-BERT's predictions. Figure 2 and Figure 3 plot the distribution of predicted TA probabilities ("*TA Adoption Probability*") across investor groups, segmented by self-declared investment approaches and by investment horizons, respectively. A desirable property of TA-BERT's predictions is that most messages fall in either the very low-probability region ($< 5\%$) or very high-probability

[13] As a benchmark, we implement a bag-of-words (BoW) approach. Specifically, we obtain a TA-related keyword list from Cookson and Niessner (2020) and compute relevance scores using TF-IDF on the same 20,000 messages. The correlation between the TF-IDF scores and GPT's classifications is only 28%. To further assess accuracy, we randomly select 500 messages where the two approaches diverge and task a PhD research assistant with manual classifications. The research assistant is instructed to classify the messages based on their knowledge of technical analysis strategies. The assistant's classification aligns with GPT's classifications 90% of the time.

region ($> 95\%$), suggesting that its classifications are typically unambiguous.

Moreover, TA-BERT's classifications are positively correlated with investors' self-declared approaches and horizons. For instance, Figure 2 shows that users identifying as technical or momentum traders rely more heavily on technical analysis than those identifying as fundamental, value, growth, or global macro investors. However, the correlation is far from perfect: many self-declared fundamental, value/growth, or global macro investors also post TA-related messages, underscoring the malleability of strategy use—a pattern we formally analyze in Section 3.

Finally, Figure 4 presents word clouds of the most frequent unigrams and bigrams in messages classified as TA-related. Panel A shows that technical messages frequently reference charts, consistent with the findings of Jiang et al. (2023), and include canonical TA terms such as "resistance," "support," and "gap." Panel B highlights bigrams that capture trading horizons, including "short-term" and "next week," alongside other familiar TA expressions.

Classifying Other Trading Strategies We follow the same two-step approach to identify the messages that reference an investment strategy and, in particular, those containing fundamental analysis (FA). The only change required is revising the prompt provided to GPT-4 Turbo. We then fine-tune specialized BERT models to identify FA-related messages. For FA, 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-

[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":}.

As with TA, we visualize FA-BERT's classifications in Figure 5 and Figure 6, grouped by self-declared investment approaches and horizons, respectively. As in TA, we observe that FA-BERT's classification is aligned with the self-declared investment approaches. For example, Figure 2 shows that investors who self-report using fundamental, value, and growth rely more heavily on fundamental analysis (FA) in their StockTwits messages. We also present word cloud plots in Figure 7 showing the high-frequency unigrams and bigrams in FA messages.

For the overall strategy identification, 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": }.

Similarly, we classify a message as Other Strategy (OS) if it refers to an investment strategy but not to either TA or FA.

Overall, about 31% of all StockTwits messages involve explicit references to trading strategies. Within this subset, 44% relate to FA, 28% to TA, and 28% to OS.

2.2.2. Message Sentiment

For each message, we obtain sentiment from the user-declared sentiment flag (available for approximately 7% of the total messages). For the remaining messages without a self-reported sentiment label, we impute sentiment using a supervised

learning approach. Specifically, we randomly select 100,000 messages with a user-provided bullish or bearish label and fine-tune a BERT model on this sample. The classifier outputs a probabilistic prediction of whether a given message is bullish.[15]

To validate the model, we apply the classifier to messages with a sentiment flag that were not included in the training sample. The classifier achieves an F1 score of 0.9, indicating high accuracy.

Our choice of BERT for sentiment classification ensures consistency across both strategy and sentiment tasks. As a robustness check, we also implement the maximum entropy approach of Cookson and Niessner (2020) to impute missing sentiment at the message level. The results, reported in Appendix Table A.6, show that our main findings remain both quantitatively and qualitatively unchanged.

We then aggregate sentiments at the firm-strategy level. Following Cookson et al. (2024a), sentiment is defined as the normalized difference between bullish and bearish messages:

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

Panel B of Table 1 reports the summary statistics of these sentiment measures, and Panel C report pairwise correlations. We find that the standard deviations are comparable across all four strategies (TA, FA, OS, and NS). The correlations across these sentiment measures are positive but modest in magnitude, ranging from 0.084 to 0.123, suggesting that different investment approaches often yield divergent assessments of the same stock.

[15] We do not employ GPT models for sentiment classification, as their training on vast corpora that may include future financial content could introduce significant look-ahead bias (see, e.g., Sarkar and Vafa, 2024). Instead, we use the BERT model, which is pretrained on a combination of book texts and English Wikipedia. Because neither financial news nor social media content is included in this pretraining, BERT is unlikely to introduce look-ahead bias in our setting (Google BERT team, 2018).

2.3. *Other Variables*

Our analysis also utilizes a set of firm-specific indicators. Trading activity is captured through *OIB* (Order Imbalance), a measure reflecting the retail marketable volume imbalance for a stock in a given week, following established methodologies from Boehmer et al. (2021) (BJZZ) and Barber et al. (2023a) (BHJOS). To further identify aggressive retail share purchase behavior, *RH Herding* is an indicator variable for stocks experiencing the top 1% of positive Robinhood user change ratio in a week, provided they had at least 100 users at the end of the prior week (Barber et al., 2022).

Our key control variables include *MAX*, the maximum one-day return in the preceding month to capture lottery-like return patterns (e.g., Bali et al., 2011). *Abnormal Turnover* quantifies unusual trading volume as the log difference between current turnover and the average of the prior four periods ($\log(1+Turnover_{i,t}) - \log(1+\frac{1}{4}\sum_{h=1}^{4} Turnover_{i,t-h})$), and *Abnormal News*, which captures unexpected change in the volume of news articles reported by Ravenpack, calculated using the same log-difference approach relative to recent historical averages ($\log(\#News_{i,t}) - \log(1+\frac{1}{4}\sum_{h=1}^{4} \#News_{i,t-h})$). Finally, we consider firm characteristics. These include *Market Capitalization* (calculated as price times shares outstanding), *Book-to-Market*, *Asset Growth*, and *Gross Profits-to-Asset*. The information environment and institutional ownership are captured by *Analyst Coverage*, representing the number of IBES equity analysts, and *Institutional Ownership*, calculated as the fraction of shares held by 13F institutional investors. Following Cookson et al. (2024a), we also create StockTwits investor Attention, which is the number of messages on a firm scaled by the total number of messages on that day.

Table A.2 details the variable construction procedure. Panel B in Table 1 reports summary statistics in the firm-day sample. Notably, the median firm capitalization is 1.2 billion, which is comparable with Giannini et al. (2018). This number indicates that our sample tilts toward larger firms, which suggests the relevance of our results is not limited to small firms with limited liquidity.

3. Retail Trading Strategy Choices

In this section, we examine the investor and message attributes linked to specific strategy usage revealed from individual messages. We find that while strategy choices correlate with self-declared approaches, horizons, and message attributes, much variation remains even after controlling for stock×investor fixed effects. We then turn to the drivers of this flexibility in retail strategy usage, focusing on the roles of public news releases, investor strategy performance in the recent period, and social feedback on strategy adoption.

3.1. Strategy Choice Across Investor and Stock Attributes

As shown in Figures 2–6, the TA-BERT's strategy classification broadly corresponds to investors' self-declared approaches and horizons, but it also reveals substantial flexibility. For instance, many investors who self-identify as technical traders often post messages referencing fundamental signals, and vice versa. This flexibility is further illustrated in Figure 1.

We formally assess the extent to which investor and message attributes explain strategy usage by estimating a series of panel regressions at the message level, focusing on those messages from users with available biographical information.[16] Our regression is specified as follows:

$$\boldsymbol{Usage}^{type}_{i,j,t,n} = \beta_1 \boldsymbol{X_j}^{investor} + \beta_2 \boldsymbol{Z_{j,i,t,n}}^{message} + \boldsymbol{FE} + \epsilon_{i,j,t,n}, \quad type \in \{TA, FA, OS\},$$

where $Usage^{type}_{i,j,t,n}$ is an indicator variable set to one if message n, posted by investor j about stock i on day t, is classified by our fine-tuned BERT models as referring to one of the strategy categories: Technical Analysis (TA), Fundamental Analysis (FA), or Other Strategy (OS).

$\boldsymbol{X_j}^{investor}$ represents a vector of indicators for investors' self-reported biographical characteristics: investment approach (e.g., *Technical Investor* equals one if an

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

investor's self-reported approach is either *Technical* or *Momentum*), investment horizon (e.g., *Long-Term Investor*), and experience level (e.g., *Novice* or *Professional*). $\boldsymbol{Z_{j,i,t,n}}^{message}$ represents a vector of message-specific attributes, including message length and the frequency of TA or FA-related terms based on the dictionary of keywords provided by Cookson and Niessner (2020) and a bag-of-words (BoW) method with TF-IDF weighting scheme to emphasize terms distinctive to certain messages and down-weight common terms. Finally, the regression specifications include an extensive set of fixed effects ($\boldsymbol{FE}$) at different levels – date, investor, stock, investor×stock – to account for unobserved heterogeneity and quantify the share of variations in strategy usage attributable to factors beyond observable attributes.

Panel A of Table 2 focuses on retail TA usage ($Usage^{TA}$). In Column (1), the explanatory variables of our primary interests are investors' self-reported investment approaches. The regression results show that self-reported approaches are significantly associated with TA adoption at the individual message level. For instance, self-declared technical investors (i.e., *Technical* or *Momentum*) are 10.3 percentage points (73.5% increase compared to the sample mean) more likely to mention TA in their messages relative to those self-identifying as other approaches (i.e., *Fundamental*, *Value*, or *Growth*). However, the R^2 is only 2.2%, indicating substantial heterogeneity unexplained by users' static self-reported profiles, due to either misreported approach or genuine within-investor variation in strategy adoption.

Column (2) augments the specification with investment horizons and investor experience. We find that short-term investors (i.e., *Day Trader* or *Swing Trader*) are significantly more likely to mention TA in their messages relative to intermediate-horizon focused investors (i.e., *Position Trader*), while long-term investors are less likely to do so. Moreover, self-reported *Professionals* are 4.7 percentage points (33.6% relative to the sample mean) more likely to use TA than the benchmark group (i.e., *Intermediate*), whereas *Novices* are significantly less likely. These results suggest that investment horizons and familiarity with technical approaches drive systematic differences, with seasoned investors more explicit in articulating

technical reasoning.

Column (3) further controls for message-specific attributes and adds date and stock fixed effects, accounting for unobserved heterogeneity in TA adoption across stocks and time. For example, investors may rely more heavily on technical analysis when analyzing emerging technology firms with limited credible information on fundamentals such as earnings or dividends. The prevalence of TA adoption may also reflect platform-level changes, such as the expansion of message-length limits in 2019 (Cookson et al., 2024a). We find that the indicator of LLM-classified TA usage is positively associated with the technical TF-IDF score and negatively associated with the fundamental TF-IDF score, suggesting that the LLM-based classification is aligned with the BoW approach. However, these traditional BoW measures, together with the stock and date fixed effects, only moderately raise the R^2 from 3.0% to 8.4%, implying that stock- and time-specific factors are not the primary drivers of TA adoption.

Finally, Column (4) adds stock$\times$investor fixed effects, controlling for all time-invariant variation in how investors tailor strategy usage across stocks. This specification increases the R^2 to 28.7%. Nevertheless, more than 70% of the variation remains unexplained, highlighting the substantial within-pair malleability in strategy adoption.

We extend our analysis to FA usage ($Usage^{FA}$) in Panel B. We find that self-reported technical investors are less likely to adopt fundamental analysis, whereas investors self-identifying with fundamental-related approaches exhibit significantly higher FA usage. Long-term, professional investors are more inclined to discuss fundamental factors in their posts. The inclusion of investor$\times$stock fixed effects increases the R^2 by about 16%, suggesting that investor-specific preferences are an important driver of FA usage. However, as with TA usage, the R^2 is still only around 32.5%, so a large share of variation (approximately 67.5%) remains unexplained.

In Panel C, we examine the usage of other strategies ($Usage^{OS}$), reflecting investment approaches unrelated to TA or FA as identified by our fine-tuned OS-BERT

model. 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 terms.

These results demonstrate that the message-level strategy classifications are broadly consistent with investors' self-declared approaches. Equally importantly, they reveal that investors adapt their strategy usage fluidly across time and stocks, underscoring the dynamic nature of retail investors' decision-making.

3.2. Strategy Malleability

Having established the variability and dynamics of investors' strategy choices, we next analyze the factors driving investors' adoption decisions and their switches across strategies. We address three key questions.

First, how does the firm-specific information flows shape investors' approach? For example, when there is greater news coverage on fundamentals, such as earnings releases, analyst reports, or other value-relevant disclosures, do investors shift toward fundamental analysis and away from technical or other strategies?

Second, how do investors respond to their own past performance using a specific strategy? Prior studies show that investors learn from their experience and often overweight personal outcomes (Kaustia and Knüpfer, 2008; Seru et al., 2010). They also tend to extrapolate recent returns (see, e.g., Greenwood and Shleifer 2014, Da et al. (2021)), engage in positive feedback trading (DeLong et al., 1990), or become more confident when outcomes are favorable (Daniel et al., 1998). Hence it remains an open question whether retail investors adjust their strategy choices in response to the performance of their own strategies.

Third, to what extent does social feedback influence retail investors' strategy adoption? Positive peer feedback (e.g., likes) to an investor when using a strategy could reinforce continued use of the same strategy, while strong validation of alternative strategies might draw the investor away. As conceptualized in Hirshleifer

(2020), Akçay and Hirshleifer (2021), and amony others, folk models generating a greater excitement ("buzz") are more likely to spread, leading more agents to adhere to them.[17]

To answer these questions, we extend the earlier analyses to examine various factors potentially influencing retail strategy usage: (i) public news releases, (ii) investor strategy performance, and (iii) social feedback on strategy adoption. Focusing on investors who have posted at least one strategy-related message in the past three months, we estimate the following panel regression model:

$$\begin{aligned} \textit{Usage}^{type}_{i,j,t,n} = \beta_1 \textit{Public News Releases}_{i,t} + \beta_2 \textit{Investor Strategy Performance}^{type}_{j,t-1} \\ + \beta_3 \textit{Social Feedback}^{type}_{j,t-1} + \beta_4 \boldsymbol{X}_{i,j,t,n} + \boldsymbol{FE} + \epsilon_{i,j,t,n}, \quad type \in \{TA, FA, OS\}. \end{aligned}$$

Public News Releases$_{i,t}$ represents a vector of three indicator variables that equal one if a specific type of news about stock i has been released on day t, including earnings related news, analyst news (e.g., recommendations, price targets), and business news (e.g., mergers, credit ratings, labor issues).

Investor Strategy Performance$_{j,t-1}$ captures strategy-specific performance over the past three months for a given investor. To compute the past performance metrics for investor j, we focus on a five-day window (i.e., $+1$ to $+5$) relative to the message's posting date, and calculate the sentiment-weighted average of the 5-day compounded DGTW excess returns across all messages, by strategy types, posted by investor j in the prior three months. To mitigate look-ahead bias, the three-month period for this calculation ends seven days before the current message date. Finally, investors are classified into high or low groups based on the 75th and 25th percentile cutoffs of their investment strategy-specific performance.

Similarly, *Social Feedback*$_{j,t-1}$ is measured by the total number of likes received across all messages, by strategy types, for investor j over the prior three months, with high and low classifications determined using the same percentile thresholds.

[17] Emerging evidence indicates that social feedback encourages greater content production on social media (Srinivasan, 2023), but little is known about whether it also alters the types of investment strategies investors employ.

The control vector X includes the frequency of usage in the three strategies over the past three months. We also include investor, stock, and date fixed effects.

We report our findings in Table 3. Panel A focuses on the relationship between news variables and strategy usage. We find that all three news indicators are positively associated with fundamental analysis (FA) usage but negatively associated with technical analysis (TA) and other strategies (OS). For example, the arrival of earnings news decreases TA usage by 2.5 percentage points (17.8% of the sample mean) and OS usage by 1.1 percentage points (10% of the sample mean), while increasing FA usage by 4.8 percentage points (roughly 28.2% of the sample mean). These results support our hypothesis that firm-specific news facilitates information processing related to fundamentals, leading to an increase in FA usage and a decline in TA and OS usage.

Panel B incorporates determinants based on an investor's past performance for each strategy type. In column (1), we analyze the determinants of TA usage. We find that investors experiencing poor TA performance in the prior quarter reduce their TA usage by 0.7 percentage points in the subsequent quarter. Conversely, TA usage increases by 0.5 percentage points when investors experience poor FA performance or high OS performance.

In column (2), we examine FA usage. Similar to TA, FA usage decreases following poor FA performance but increases after high TA performance. Unexpectedly, high OS performance is positively related to FA usage, contrary to our expectation that it would reduce FA usage. In column (3), we analyze OS usage and observe that poor FA performance is positively associated with increased OS usage. Overall, these findings suggest that investors dynamically adjust their strategy choices based on their past performance, generally moving away from underperforming strategies and adopting alternatives with relatively better outcomes.

Panel C examines the role of social feedback on strategy usage. We introduce the total number of likes received for each strategy type as a measure of social feedback, alongside controls for past strategy usage, performance determinants,

and other variables. In column (1), where the dependent variable is TA usage, we find that positive social feedback for TA messages increases future TA usage by 1.2 percentage points, whereas negative social feedback decreases TA usage by 1.4 percentage points. Furthermore, TA usage increases when investors receive negative feedback for FA or OS-related messages and decreases when the social feedback from FA or OS messages is positive.

Columns (2) and (3) explore the effects of social feedback on FA and OS usage, respectively, revealing a similar pattern: investors increase strategy usage in response to positive feedback on that strategy and reduce usage when other strategies receive favorable feedback.

Importantly, the coefficients for social feedback are generally larger and more significant than those for past performance. For example, investors who used TA in the previous quarter but performed poorly are 0.7 percentage points less likely to use TA in the following quarter. By contrast, conditional on performance, those who posted TA-related messages but received few likes are 1.3 percentage points less likely to continue using TA—an effect that is 85% larger than the response to their own performance.

These findings are consistent with experimental evidence showing that social feedback motivates greater content production on social media Srinivasan (2023) and underscore the importance of incorporating social dimensions into the study of investment behavior (e.g., Hirshleifer, 2020; Han et al., 2022; Cookson et al., 2024b).

In summary, our results demonstrate that while investors may initially self-identify with a given investment strategy, they adjust their approach based on firm-specific news, their own past performance, the past performance of alternative strategies, and social feedback from peers. These findings provide novel insights into how both personal performance attributes and social cues shape investor learning and belief formation on social networks.

4. Performances of Retail Investment Strategies

A large literature examines retail investors' trading activities and performance, typically focusing either on aggregate market order flows or on individual trading accounts (e.g., Barber and Odean, 2013; Boehmer et al., 2021; Barber et al., 2023a). Much less is known, however, about the heterogeneity in strategies retail investors pursue. In particular, we lack systematic evidence on how investors adopt specific strategies and how these choices translate into investment performance. In this section, we fill this gap by evaluating the performance of distinct retail investment strategies, as identified through our LLM-based classification of StockTwits messages.

4.1. Daily Return Predictability

We first test whether strategy-specific sentiment on StockTwits predicts next-day stock returns. Following Cookson et al. (2024a), we estimate the stock-day predictive regression:

$$\textit{Return}_{i,t+1} = \beta_1 \textit{Sentiment}_{i,t}^{type} + \beta_2 \textit{Attention}_{i,t} + \gamma \boldsymbol{X}_{i,t} + \boldsymbol{\delta}_t + \epsilon_{i,t+1},$$

where $\textit{Return}_{i,t+1}$ denotes stock i's return on the next trading day $t+1$. $\textit{Sentiment}_{i,t}^{type}$ (with $type \in \{\text{TA, FA, OS, NS}\}$) measures sentiment extracted from messages classified by LLMs into four categories: Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), and Non-Strategy (NS). $\textit{Attention}_{i,t}$ captures StockTwits investor attention to stock i on day t, defined as the fraction of StockTwits messages about stock i relative to the total number of messages across all stocks that day. The control vector $\boldsymbol{X}_{i,t}$ includes the logarithm of market capitalization, the logarithm of book-to-market, asset growth, gross profitability, analyst coverage, the logarithm of institutional ownership, the maximum daily return (MAX) in the previous month, abnormal turnover, abnormal news volume, and returns at five daily lags. All regressions include trading-day fixed effects, and standard errors

are clustered by trading day.

The results reported in Table 4 reveal striking heterogeneity across strategy types. Columns (1)–(4) consider strategy-specific sentiment measures separately. TA sentiment significantly and *negatively* predicts returns: stocks with more bullish TA messages underperform on the following day. Specifically, stocks with the most bearish TA sentiment ($Sentiment^{TA} = -1$) outperform those with the most bullish sentiment ($Sentiment^{TA} = 1$) by about 3.2 bps (1.6×2) the next day.

In contrast, FA sentiment significantly and *positively* predicts returns: stocks with the most bullish FA sentiment outperform those with the most bearish by 2.8 bps (1.4×2).

OS sentiment also negatively predicts returns, similar to TA. In contrast, NS sentiment is uninformative. Column (5) reports a horse-race regression including all four sentiment measures simultaneously. The coefficients remain stable, reflecting their low correlations and underscoring the robustness of these findings.[18]

These results provide direct evidence of heterogeneity in the informativeness of retail strategies. The negative return predictability of TA and OS sentiments is consistent with the view that retail investors often engage in attention-driven or speculative trading, generating short-term mispricing and subsequent reversals (e.g., Barber and Odean 2000; Kumar and Lee 2006; Barber et al. 2022; Bryzgalova et al. 2023). By contrast, the positive predictive power of FA sentiment suggests that some retail investors incorporate valuable fundamental information into their analyses, consistent with the "informed retail trader" perspective (e.g., Kaniel et al. 2008; Kelley and Tetlock 2013; Boehmer et al. 2021; Welch 2022). Overall, these findings reconcile prior literature and underscore that retail trading cannot be characterized as uniformly "noise" or "smart money," but instead depends critically on the type of strategy emphasized.

[18]Table A.7 in the Appendix presents subsample tests, showing that our findings on the return predictability are robust, especially during the recent periods when retail trading is active, and particularly stable before and after the GME short-squeeze episode in 2021.

4.2. *Long-Short Strategies*

To quantify the economic magnitude of the return predictability of strategy sentiment, we construct long-short (L/S) portfolios separately for TA, FA, and OS sentiments. Following the signal-based strategy construction methodology in Jensen et al. (2023), the L/S strategy takes positions across the entire cross-section of stocks with valid sentiment measures.[19] To mitigate confounding factors, we residualize sentiment by employing the approach of Nagel (2005). Specifically, on each day, we first estimate a cross-sectional regression of daily sentiment measures on a set of daily stock characteristics – investor attention, the logarithm of market capitalization, abnormal turnover, and returns at five lags – and then use the residual sentiments to form the L/S trading strategies. For ease of interpretation, TA and OS sentiment scores are multiplied by -1, reflecting their negative relation to returns.

Panel A of Table 5 reports annualized average daily returns and Sharpe ratios. Consistent with the predictive regression results in Table 4, the strategy that bets against retail TA sentiment (long stocks with low TA sentiment, short stocks with high TA sentiment) is profitable, generating an annualized return of 9.5% (t-statistic = 2.91) and a Sharpe ratio of 0.86. In contrast, the FA-based L/S strategy delivers a significant annual return of 6.48% (t-statistic = 2.04), with a Sharpe ratio of 0.58. These results are robust, and even stronger, to DGTW performance adjustments (Daniel et al., 1997), as shown in Panel B, where the TA (FA) L/S strategy yields an annualized DGTW excess return of 10.10% (7.75%) and a Sharpe ratio of 1.00 (0.75).

These results highlight that the economic value of retail sentiment depends

[19] Specifically, the L/S strategy weights are determined proportionally to each stock's deviation from the cross-sectional average sentiment, calculated as:

$$r_t^{L/S,type} = \frac{\sum_{j=1}^{N}(S_{j,t-1}^{type} - S_{t-1}^{type})r_{j,t}}{\frac{1}{2}\sum_{j=1}^{N}\left|S_{j,t-1}^{type} - S_{t-1}^{type}\right|}, \quad \text{where } S_{t-1}^{type} = \frac{1}{N}\sum_{j=1}^{N} S_{j,t-1}^{type}, \quad type \in \{TA, FA, OS\}.$$

.

critically on the underlying strategy. Contrarian signals extracted from TA and OS sentiment generate sizable abnormal returns, whereas FA sentiment aligns with informed trading and yields positive returns. Although such strategies may not be implementable in practice due to high turnover and transaction costs, the estimated returns quantify the role of retail sentiment in shaping return dynamics and market frictions.

4.3. Return Predictability at Longer Horizons

We next examine the return predictability of strategy-specific sentiments over longer horizons. To test whether our findings extend beyond the very short run, we re-estimate the predictive regression using cumulative stock returns over three windows: $t+1$ to $t+5$, $t+6$ to $t+10$, and $t+11$ to $t+15$. The results are reported in Table 6.

The return predictability documented in Table 4 is strongest in the one-week window ($t+1$ to $t+5$) for all three strategy-related sentiments (TA, FA, and OS). Importantly, TA sentiment also significantly predicts returns at the intermediate horizon ($t+6$ to $t+10$), while OS sentiment's predictive power persists even further, extending to the longest horizon ($t+11$ to $t+15$). By contrast, FA sentiment significantly predicts positive returns over the one-week horizon but loses significance at longer horizons, with no evidence of subsequent reversal.[20]

These findings reinforce our one-day evidence and underscore sharp differences in informativeness across strategies. FA sentiment reliably predicts near-term positive returns, consistent with the timely incorporation of fundamentals. By contrast, TA and OS sentiment exhibit persistent negative predictability, extending well beyond the short run and suggesting that these signals reflect price movements driven by persistent behavioral biases rather than value-relevant information. In-

[20]We conduct a series of robustness checks. Table A.3 in the Appendix estimates predictive regressions using the Fama–MacBeth approach. Table A.4 repeats the analysis with DGTW-adjusted returns. Table A.5 imposes a minimum of 10 messages per firm-day. Table A.6 uses sentiment inferred via the algorithm of Cookson and Niessner (2020) instead of BERT. All tests yield qualitatively similar results.

vestors who follow such sentiment are thus more likely to incur losses, whereas those emphasizing fundamentals capture genuine informational advantages.

4.4. Return Predictability and Investor Sophistication

Investment performance may vary significantly with investor sophistication. While our evidence shows that retail TA sentiment is generally on the wrong side of the market, prior research documents technical indicators that consistently generate abnormal returns (e.g., Jegadeesh and Titman, 1993b; Han et al., 2013, 2016; Jiang et al., 2023; Murray et al., 2024). This contrast suggests that more sophisticated investors on StockTwits may be able to better discern valuable technical signals, making sentiment more informative on days when they are more active.

To examine this possibility, we measure message sophistication (*FracProMsg*) as the fraction of messages contributed by self-reported professional investors for each stock-day. We then interact our strategy-specific sentiment measures with *FracProMsg* and re-estimate the predictive regression, including both the interaction term and the standalone *FracProMsg*. The corresponding regression results are reported in Table 7.

We find the interaction between *FracProMsg* and TA or OS sentiment is positive and significant, attenuating the otherwise negative return predictability of these sentiment measures. By contrast, the predictive power of FA sentiment remains unaffected by professional participation, consistent with the notion of "*Wisdom of Clouds*" that fundamental signals are informative even when discussions are dominated by novice traders (see, e.g., Chen et al., 2014; Welch, 2022).[21]

In summary, these findings suggest that greater participation by professional users dampens the irrational sentiment embedded in technical and speculative discussions, whereas fundamental analysis sentiment remains robust.

[21]The standalone coefficient on *FracProMsg* is positive and significant in predicting one-week-ahead returns, with no evidence of reversal in the subsequent two weeks. Therefore, the presence of sophisticated users in StockTwits discussions—regardless of whether their sentiment is bullish or bearish—appears to predict higher future returns. One interpretation is that informed traders are more likely to post when they observe bullish signals.

5. Retail Strategies, Trading, and Informativeness

Our analysis thus far shows that StockTwits investor sentiments associated with different investment strategies yield markedly different investment outcomes. A natural question next is whether investors merely "talk the talk" without actually "walking the walk." In other words, do social media discussions reflect genuine trading behavior, or are they simply cheap talk? Because StockTwits does not provide brokerage-level trading data, we cannot directly observe whether investors' trades align with their posts. Instead, we assess whether StockTwits sentiment is representative of retail investor trading by examining its relationship with aggregate retail market order flows, the informativeness of these flows, and the activities of Robinhood traders.

5.1. Aggregate Retail Market Order Imbalance

Following Boehmer et al. (2021) and Barber et al. (2022), we identify retail market orders and construct two alternative measures of retail order imbalance (*OIB*): $OIB_{i,t}^{BJZZ}$ and $OIB_{i,t}^{BHJOS}$. We then estimate the following panel regression at the stock-day level:

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

To align closely with market orders, the sentiment measures are constructed using only messages posted during regular trading hours (9:30–16:00).

Panel A of Table 8 reports the regression results. Column (1) shows that sentiments across all four strategy categories are positively associated with OIB^{BJZZ}. A one–standard deviation increase in TA (FA) sentiment corresponds to a 0.39 (0.42) percentage point increase in retail order imbalance. Column (2) confirms these results using OIB^{BHJOS}, again showing strong positive contemporaneous relationships. Overall, these findings demonstrate that intraday StockTwits sentiment closely tracks retail trading during market hours, providing credible information

about investors' actual trades.[22]

5.2. *Retail Order Flow Informativeness*

Boehmer et al. (2021) show that aggregate retail market order imbalances are informative and positively predict future stock returns. Building on this result, we decompose retail order flows into strategy-specific components to assess how different strategy types contribute to, or detract from, informativeness.

Specifically, we first regress *OIB* on strategy-specific sentiments using the specification in the previous section. The fitted values, $OIB^{TA}_{i,t}$, $OIB^{FA}_{i,t}$, $OIB^{OS}_{i,t}$, and $OIB^{NS}_{i,t}$, defined as $\widehat{\beta}^{type} \times Sentiment^{type}_{i,t}$, represent the components of *OIB* attributable to each sentiment type. The residual from this regression, $OIB^{Resid}_{i,t}$, captures the portion of retail order flows uncorrelated with StockTwits sentiments, which may reflect either information withheld from social media or pure noise.

We then regress next-day stock returns on the decomposed components:

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

The regression results are reported in Panel B of Table 8. Column (1) uses the decomposition based on OIB^{BJZZ}, while Column (2) employs the alternative OIB^{BHJOS} measure.

Several findings stand out. First, TA- and OS-driven order imbalance *negatively* predict future stock returns, indicating that these strategies reduce the informativeness of retail orders. Second, FA-driven order imbalance positively predict returns, enhancing informativeness. Third, NS-driven order imbalance exhibit little predictive power. Finally, the residual component positively and significantly predicts stock returns, suggesting that retail order flows contain valuable information beyond what is shared on StockTwits.[23]

[22] In a placebo test, we show that *OIB* is not significantly associated with sentiments of messages posted during the prior overnight period.

[23] Table A.8 in the Appendix reports results for longer horizons ($t+1$ to $t+5$, $t+6$ to $t+10$, and $t+11$ to $t+15$). Table A.10 presents rolling-window regressions to decompose *OIB*. In both cases, we find our findings still hold well.

Regarding economic magnitudes, a one-standard-deviation increase in FA-driven OIB leads to a 1.1 basis-point increase in next-day returns. By comparison, a similar increase in the residual component of OIB is associated with a 4.6 basis-point increase. This implies that although many informed retail investors, as expected, refrain from disclosing their information on social media and instead trade directly on it, the portion that is revealed—about 24% of that withheld—is still economically meaningful.[24] This result also contrasts with the conventional view that informed investors remain entirely secretive and raises new questions about how investor behavior has evolved in the era of social media.

Our findings in this section shed light on the central debate over whether retail investors are informed traders or primarily driven by behavioral biases. By decomposing retail order flows into strategy-specific components, we highlight important heterogeneity: retail investors who rely on fundamental signals contribute positively to price discovery, whereas those following technical or other speculative signals submit order flows that move opposite to future price changes.

5.3. Intense Buying by Robinhood Investors

We next turn to a prominent subset of retail investors, those trading on Robinhood, a leading zero-commission platform. Many Robinhood users are first-time investors,[25] and, as documented by Barber et al. (2022), they are particularly susceptible to attention-driven factors and speculative trading. These investors often engage in intense buying episodes ("herding"), which are typically followed by sharp price reversals. Since social media plays a central role in shaping retail attention, we hypothesize that Robinhood investors are particularly responsive to online discussions. This subsection therefore investigates the extent to which StockTwits sentiment explain Robinhood herding behavior.

Following Barber et al. (2022) and Welch (2022), we employ RobinTrack data

[24]The standard deviations of strategy-specific OIB measures are reported in Table A.10 in the Appendix.

[25]For example, see https://newsroom.aboutrobinhood.com/robinhood-raises-280-million-in-series-f-funding-led-by-sequoia/.

from May 2018 to August 2020. As in Barber et al. (2022), we define episodes of intense buying—"buy herding"—as days on which stock i ranks among the top 10 stocks by the daily percentage increase in Robinhood users, conditional on at least 100 users holding the stock on day $t-1$. We then estimate the following stock-day panel regression:

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

Table 9 presents the regression results. Columns (1) through (4) separately analyze the contemporaneous relationship between Robinhood herding and StockTwits sentiment for each investment strategy category (TA, FA, OS, NS). We find that all sentiment categories are positively associated with buy herding, but the economic magnitude is substantially larger for TA sentiment. Column (5) includes all sentiment categories, and the results remain consistent: each sentiment type predicts herding, with TA sentiment showing the strongest relationship.

The economic magnitude is also highly significant. A one–standard deviation increase in TA sentiment raises the probability of herding by 0.062 percentage points (18% increase compared to the unconditional frequency of herding of 0.35 percentage points), compared to 0.038 for FA sentiment (11% increase compared to the unconditional frequency) and 0.033 (10% increase compared to the unconditional frequency) for OS sentiment.

These findings highlight a strong link between StockTwits sentiment and Robinhood investor behavior. Among the different strategy types, TA sentiment displays the most pronounced association with crowded buying episodes, consistent with the relatively poor performance of TA-based strategies. This pattern resonates with the "crowded-trade" effect described by Stein (2009). In Stein's framework, quantitative strategies can underperform when traders, even if fully rational, cannot observe in real time how many others are following the same model. This coordination failure amplifies price movements away from fundamentals.

Extending this logic to retail investors, the effect may be even more pronounced.

Technical signals are both salient and easy to grasp, making them especially appealing to inexperienced traders. However, retail investors often neglect the possibility that many others may be acting on the same signals in a similarly naïve fashion. Their limited use of shorting further tilts trades toward the buy side, generating crowded buying and sharp reversals. This provides a possible explanation for our earlier findings: TA sentiment not only relates to retail herding but also foreshadows the subsequent underperformance of TA-based strategies.

6. Conclusion

This study demonstrates the potential of integrating rich social media data with LLMs to better understand investor trading strategies, sentiment, and trading behavior. By extracting trading strategies from social media posts, we show that retail investors dynamically adjust their strategies in response to firm-specific news, past performance, and peer feedback. Technical analysis (TA) becomes more prevalent when firm-specific news is scarce, while fundamental analysis (FA) gains traction as news flow intensifies. Strategy adoption is also reinforced by prior success with the strategy and by social validation, such as likes on messages.

Our empirical analysis reveals sharp differences in strategy effectiveness. TA- and OS-based sentiment signals negatively predict returns, especially when social media discussions are dominated by less sophisticated users, whereas FA-based sentiment positively predicts returns. All sentiment categories are strongly associated with retail buying, but TA sentiment stands out: FA-linked order flows are informative about future returns, while TA- and OS-driven flows predict returns in the wrong direction. TA sentiment also exhibits the strongest connection to intense buying episodes on Robinhood, underscoring its role in crowded trading and sharp reversals.

Overall, our study illustrates how combining social media data with LLMs provides a powerful new lens on retail trading—revealing how strategies shaped by

social interaction can both inform markets and destabilize them. Looking ahead, this approach offers new opportunities for academic research, such as examining how retail behavior interacts with institutional trading or influences broader market dynamics. It also provides useful perspective for policy discussions by helping identify conditions under which retail participation supports price discovery versus when it may amplify vulnerabilities in financial markets.

REFERENCES

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

Akçay, Erol, and David Hirshleifer, 2021, Social finance as cultural evolution, transmission bias, and market dynamics, *Proceedings of the National Academy of Sciences of the United States of America* 118, e2015568118.

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.

Bali, Turan G, Nusret Cakici, and Robert F Whitelaw, 2011, Maxing Out: Stocks As Lotteries and The Cross-Section of Expected Returns, *Journal of financial economics* 99, 427–446.

Barber, Brad M, Xing Huang, Philippe Jorion, Terrance Odean, and Christopher Schwarz, 2023a, 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, Shengle Lin, and Terrance Odean, 2023b, Resolving a Paradox: Retail Trades Positively Predict Returns but Are Not Profitable, *Journal of Financial and Quantitative Analysis* 1–35.

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).

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2018, Extrapolation and bubbles, *Journal of Financial Economics* 129, 203–227.

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.

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, Shuaiyu, Lin Peng, and Dexin Zhou, 2025, Man versus machine: How AI exploits retail sentiment in financial markets, Working paper.

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, 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, 2024a, The Social Signal, *Journal of Financial Economics* 158, 103870.

Cookson, J Anthony, William Mullins, and Marina Niessner, 2024b, 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.

Da, Zhi, Xing Huang, and Lawrence J Jin, 2021, Extrapolative beliefs in the cross-section: What can we learn from the crowds?, *Journal of Financial Economics* 140, 175–196.

Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, *The Journal of finance* 52, 1035–1058.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *The Journal of Finance* 53, 1839–1885.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, Positive feedback investment strategies and destabilizing rational speculation, *Journal of Finance* 45, 375–395.

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.

Dou, Winston Wei, Itay Goldstein, and Yan Ji, 2024, AI-Powered Trading, Algorithmic Collusion, and Price Efficiency, *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper* .

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.

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* .

Google BERT team, 2018, BERT Base Uncased Model, `https://huggingface.co/google-bert/bert-base-uncased`, Hugging Face Model Card; Accessed on 26 August 2025.

Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of returns and expected returns, *The Review of Financial Studies* 27, 714–746.

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

Han, Bing, David Hirshleifer, and Johan Walden, 2022, Social transmission bias and investor behavior, *Journal of Financial and Quantitative Analysis* 57, 390–412.

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, 2020, Presidential address: Social transmission bias in economics and finance, *The Journal of Finance* 75, 1779–1831.

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.

Jensen, Theis Ingerslev, Bryan Kelly, and Lasse Heje Pedersen, 2023, Is there a replication crisis in finance?, *The Journal of Finance* 78, 2465–2518.

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.

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

Kaustia, Markku, and Samuli Knüpfer, 2008, Do investors overweight personal experience? evidence from ipo subscriptions, *The journal of finance* 63, 2679–2702.

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.

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.

Nagel, Stefan, 2005, Short sales, institutional investors and the cross-section of stock returns, *Journal of financial economics* 78, 277–309.

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.

Sarkar, Suproteem K, and Keyon Vafa, 2024, Lookahead bias in pretrained language models, *Available at SSRN 4754678* .

Seru, Amit, Tyler Shumway, and Noah Stoffman, 2010, Learning by trading, *The Review of Financial Studies* 23, 705–739.

Srinivasan, Karthik, 2023, Paying attention, Technical report, Technical Report, Mimeo.

Stein, Jeremy C, 2009, Presidential Address: Sophisticated Investors and Market Efficiency, *The Journal of Finance* 64, 1517–1548.

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

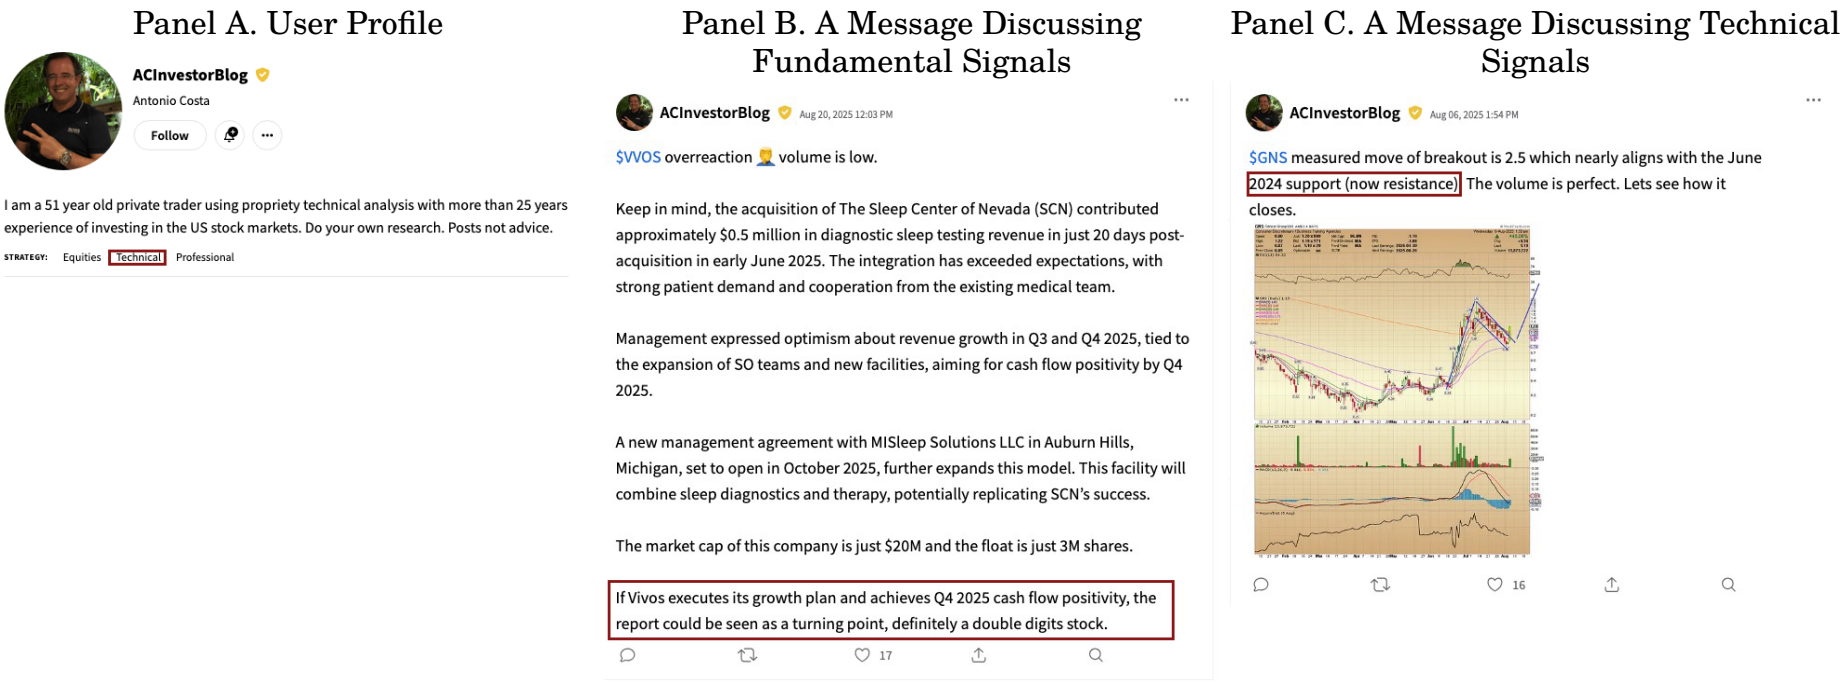


Fig. 1. **Strategy Malleability: An Illustration**

This figure presents anecdotal evidence of the malleability of investor strategies. The screenshots, sourced from the StockTwits platform, feature a user, ACInvestorBlog, who self-identifies as a "*Technical*" strategist (as indicated in the user profile in Panel A). In Panel B, the user discusses fundamental indicators in a post from August 20, 2025, stating "*If Vivos executes its growth plan and achieves Q4 2025 cash flow positivity, the report could be seen as a turning point, definitely a double digits stock.*" In contrast, Panel C shows the user discussing technical signals in a post from August 6, stating "*$GNS measured move of breakout is 2.5 which nearly aligns with the June 2024 support (now resistance) The volume is perfect. Lets see how it closes.*" The red rectangular boxes have been added for emphasis and are not part of the original screenshots.

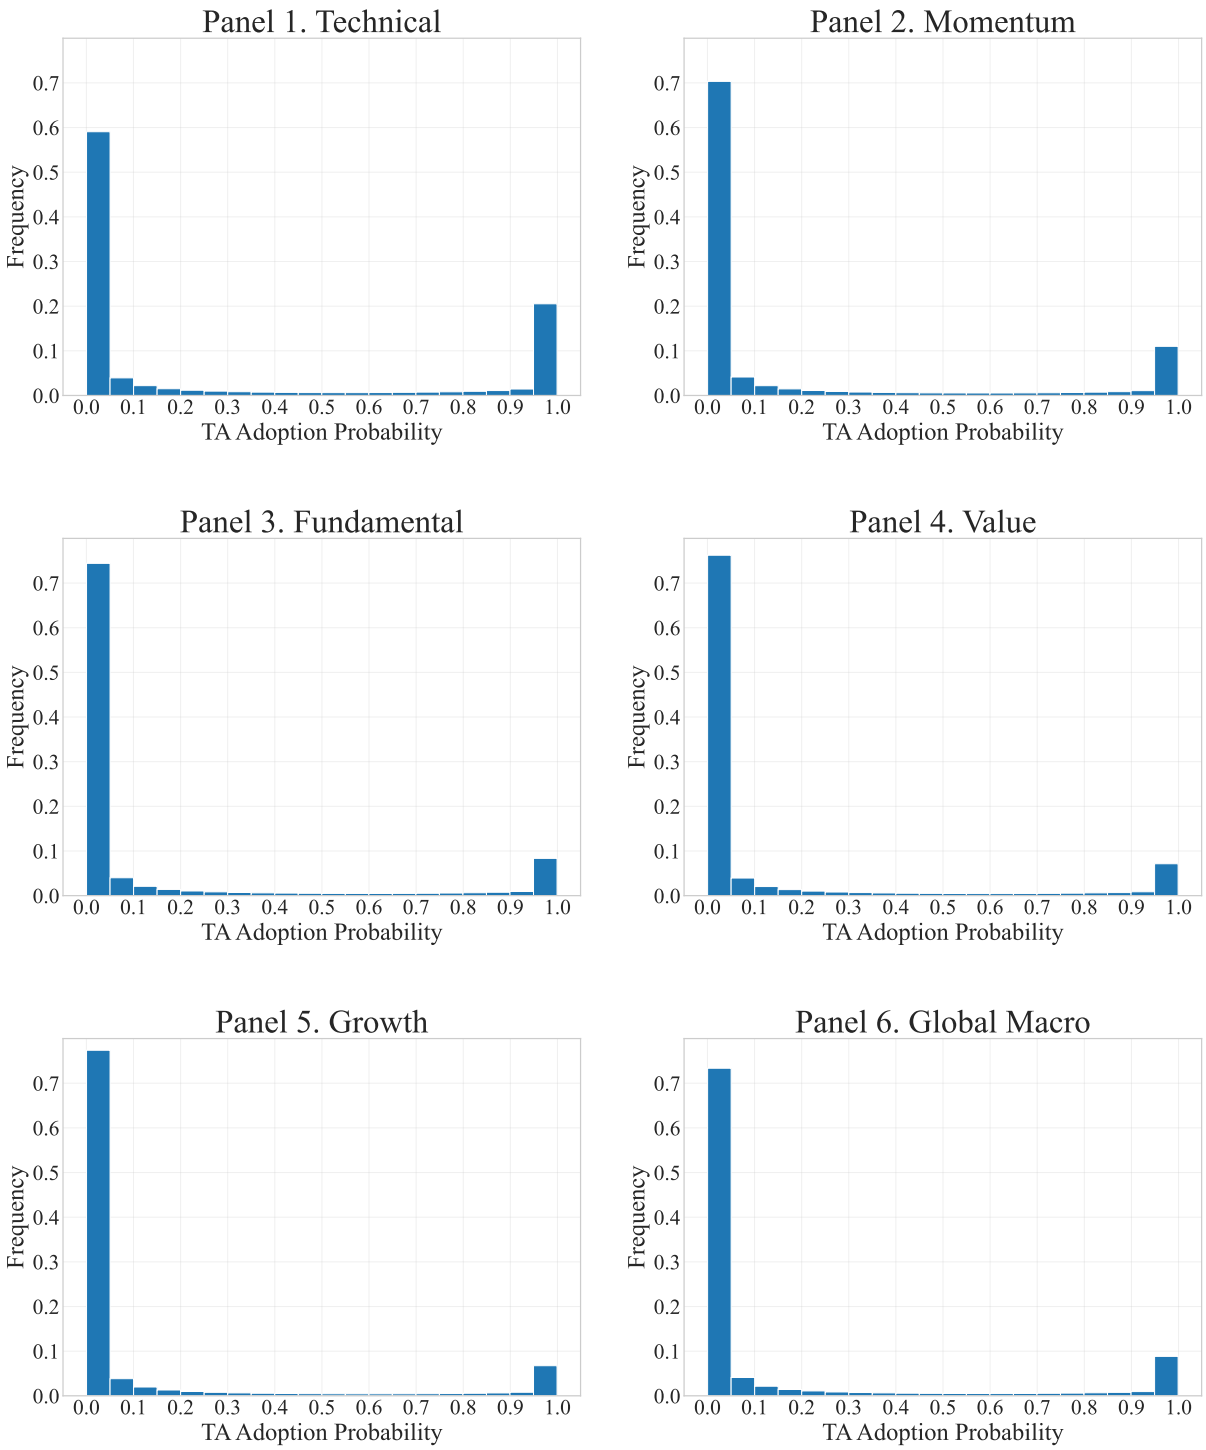


Fig. 2. **TA Adoption Probability Across Self-Declared Approaches**

This figure plots the distribution of *TA Adoption Probability* for StockTwits messages, grouped by users' self-declared investment approaches. Each message receives a probabilistic score from the *TA-BERT* model, where higher values indicate a greater likelihood that the message employs technical analysis. For each investment approach subgroup, the histograms report the frequency of messages across probability levels.

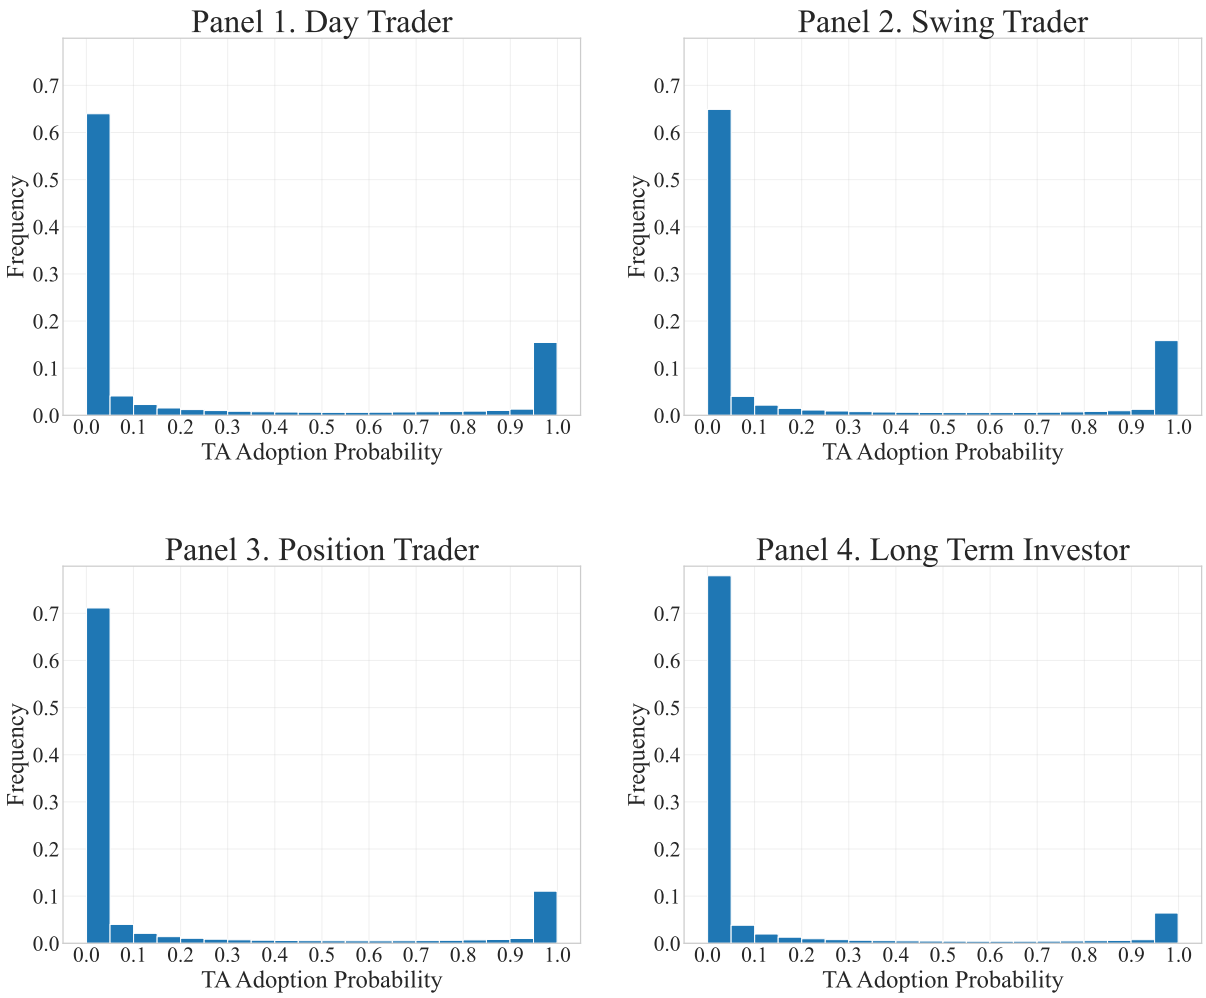


Fig. 3. **TA Adoption Probability Across Self-Declared Investment Horizons**

This figure plots the distribution of *TA Adoption Probability* for StockTwits messages, grouped by users' self-declared investment horizons. Each message receives a probabilistic score from our fine-tuned *TA-BERT* model, where higher values indicate a greater likelihood that the message employs technical analysis. For each investment horizon subgroup, the histograms report the frequency of messages across probability levels.

Panel A. Word Cloud of Unigrams in Technical Messages



Panel B. Word Cloud of Bigrams in Technical Messages

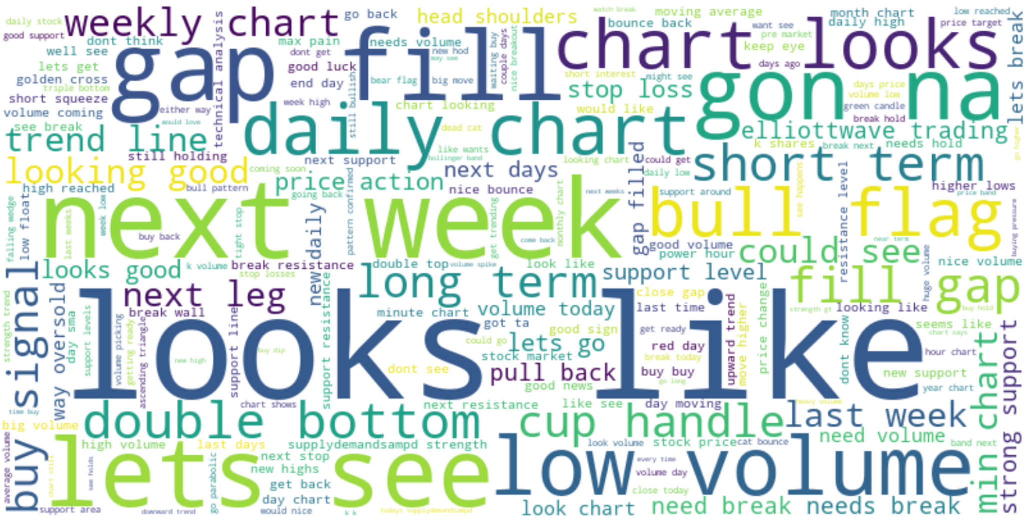


Fig. 4. **Word Clouds of Technical Analysis Messages**

This figure presents word clouds derived from StockTwits messages classified as TA-related (*TA Adoption Probability* ≥ 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.

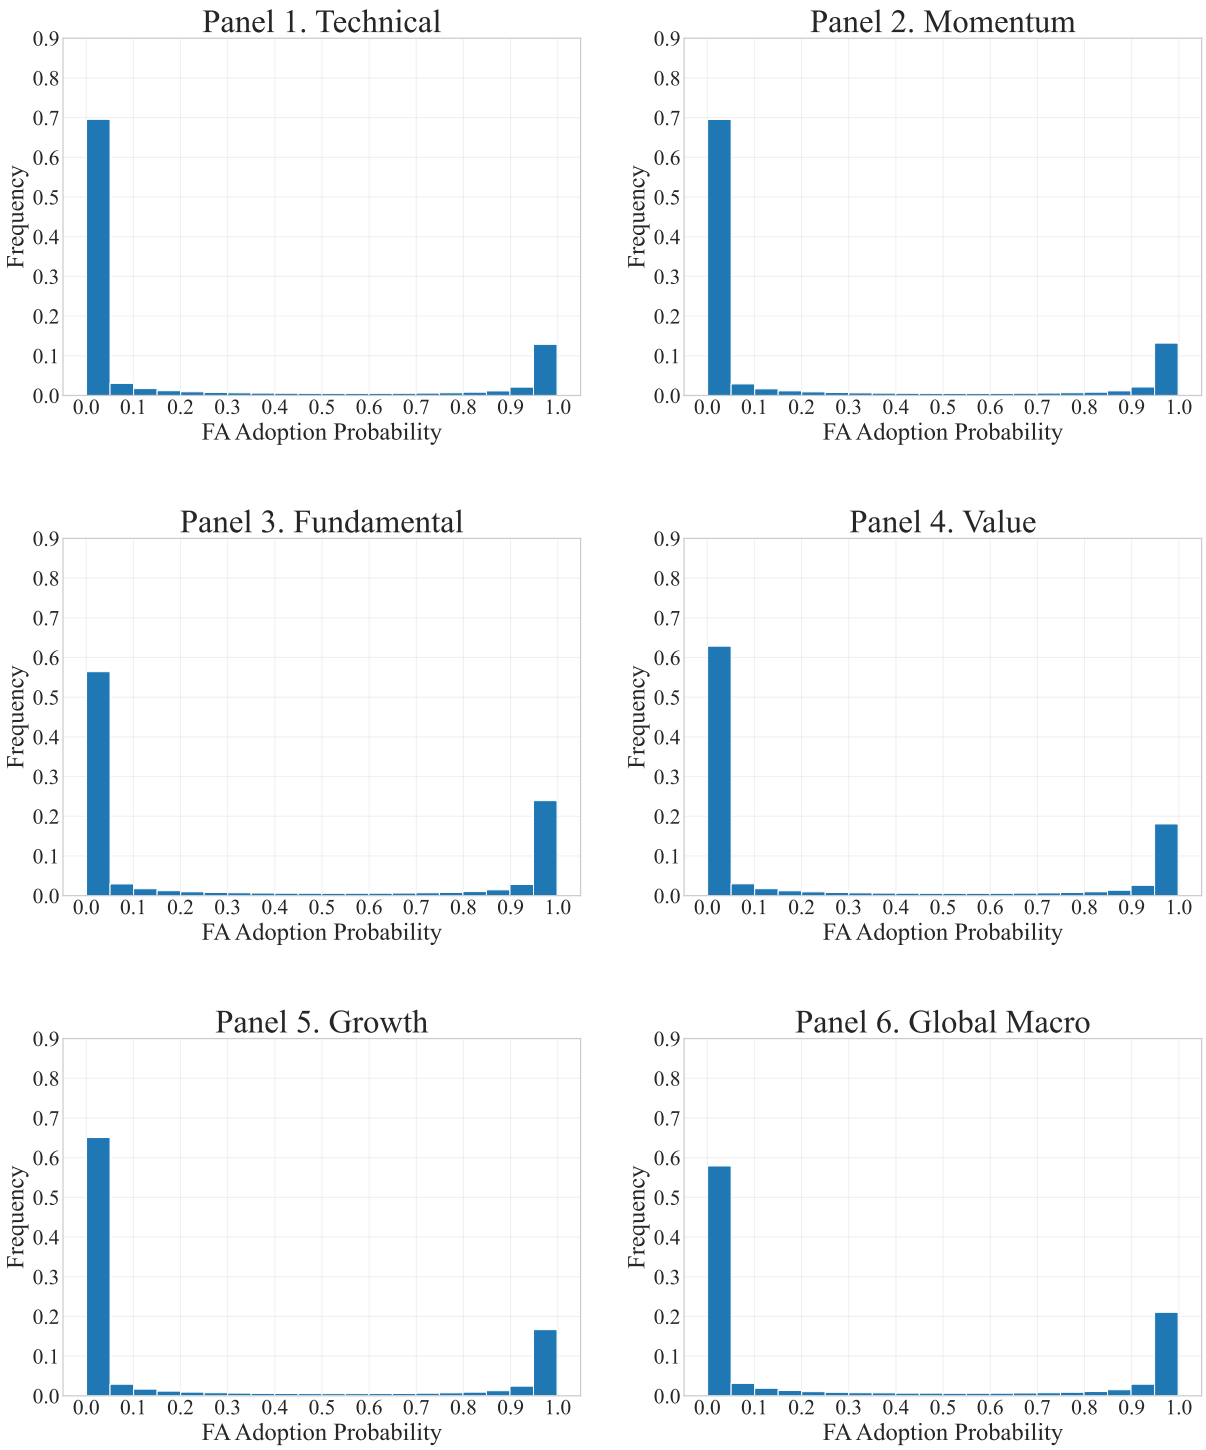


Fig. 5. **FA Adoption Probability Across Self-Declared Approaches**

This figure shows the distribution of *FA Adoption Probability* for messages posted on StockTwits, grouped by users' self-declared investment approaches. Each message receives a probabilistic score from the *FA-BERT* model, where higher values indicate a greater likelihood that the message employs technical analysis. For each investment approach subgroup, the histograms report the frequency of messages across probability levels.

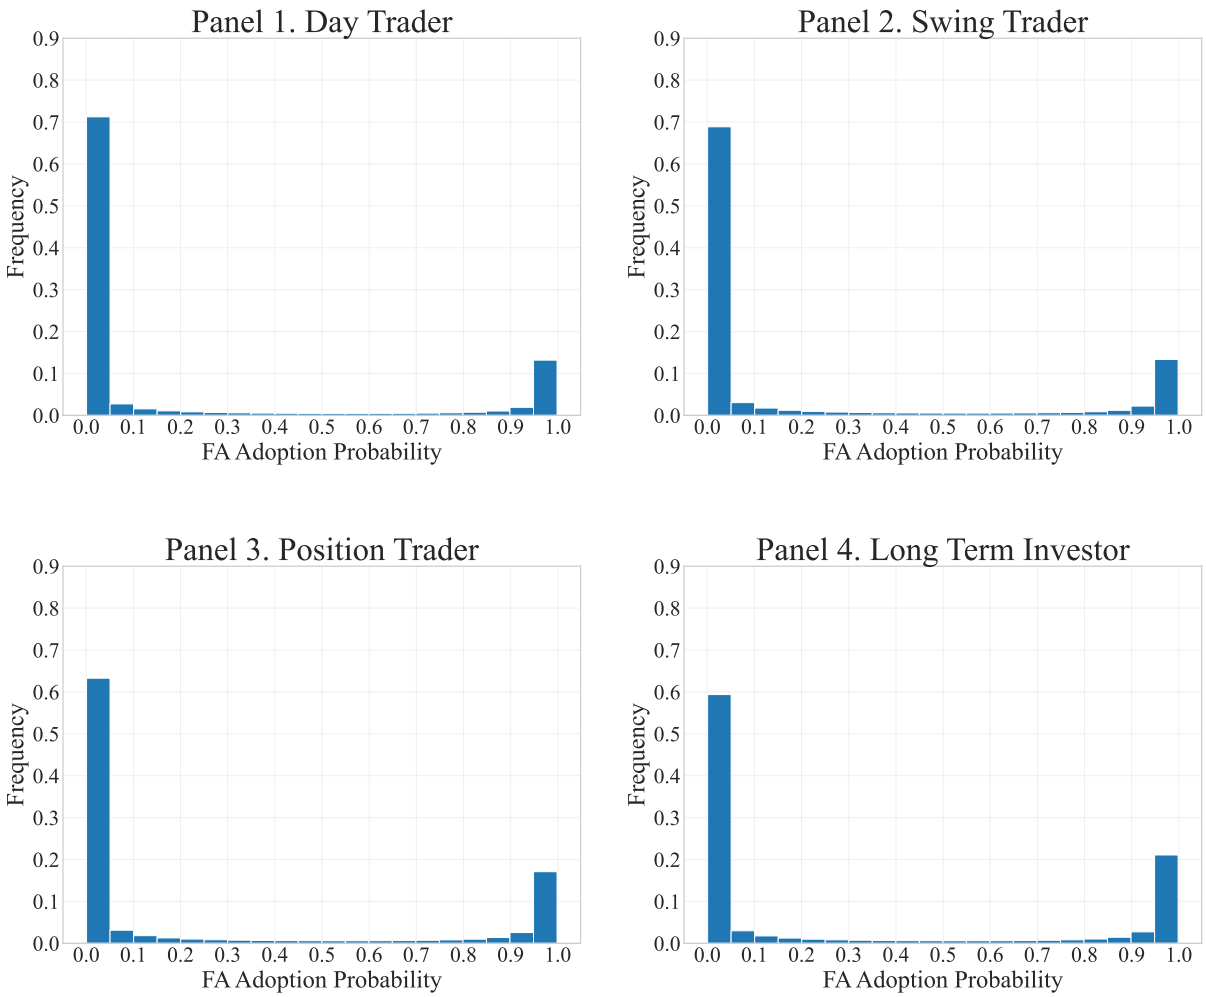


Fig. 6. **FA Adoption Probability Across Self-Declared Horizons**

This figure presents the distribution of *FA Adoption Probability* for messages posted on StockTwits, grouped by users' self-declared investment horizons. Each message receives a probabilistic score from our fine-tuned FA-BERT model, where higher values indicate a greater likelihood that the message employs technical analysis. For each investment horizon subgroup, the histograms report the frequency of messages across probability levels.

Panel A. Word Cloud of Unigrams in Fundamental Messages



Panel B. Word Cloud of Bigrams in Fundamental Messages



Fig. 7. **Word Clouds of Fundamental Analysis Messages**

This figure presents word clouds derived from StockTwits messages classified as FA-related (*FA Adoption Probability* ≥ 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 from StockTwits users with available self-reported biographical data. The reported variables include the indicator of LLM-classified trading strategies (Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS)), self-declared 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 analyses based on the 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 strategy types: TA, FA, OS, and NS (Non-Strategy). We also report StockTwits user attention (Cookson et al., 2024a), retail market order imbalance (OIB) based on methodologies from Boehmer et al. (2021) and Barber et al. (2023a), as well as firm characteristics including the logarithm of market cap, book-to-market, asset growth, gross profitability, analyst coverage, institutional ownership (IO), maximum daily return in the prior month (MAX), abnormal turnover, and abnormal news article volume. Panel C reports correlations between sentiment scores across strategy types. The sample period spans January 2010 to June 2023. Table A.2 in Appendix shows variable definitions.

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-Assets	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
IO	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 Volume	2,974,934	-0.50	-0.53	0.98	-1.71	-1.15	0.00	0.65

Panel C: Correlations Between Sentiments across Strategy Types

	Sentiment^{TA}	Sentiment^{FA}	Sentiment^{OS}	Sentiment^{NS}
Sentiment^{TA}	1.000			
Sentiment^{FA}	0.128	1.000		
Sentiment^{OS}	0.127	0.097	1.000	
Sentiment^{NS}	0.090	0.084	0.102	1.000

Table 2: **LLM-Classified Retail Strategies and Investor/Message Attributes**

This table reports results from panel regressions relating retail investors' usage of investment strategies, classified by LLMs from StockTwits messages, to the two sets of covariates: (i) self-declared investor attributes, and (ii) message-specific attributes. Panels A, B, and C correspond to the message-level usage of Technical Analysis (TA), Fundamental Analysis (FA), and Other Strategies (OS), respectively. The regression is estimated at the individual message level, with the following specification:

$$Usage_{i,j,t,n}^{type} = \beta_1 \boldsymbol{X_j}^{investor} + \beta_2 \boldsymbol{Z_{j,i,t,n}}^{message} + \boldsymbol{FE} + \epsilon_{i,j,t,n}, \quad type \in \{TA, FA, OS\}$$

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 LLMs. Investor-level characteristics ($X_j^{investor}$) include self-reported attributes such as Technical Investor, Long-Term Investor, and Professional Investor. Message-specific attributes ($Z_{j,i,t,n}^{message}$) include the TF-IDF weighted scores for technical and fundamental keywords based on the word lists from Cookson and Niessner (2020), and the logarithm of the number of words. We consider various combinations of fixed effects, including date, stock, investor, and stock×investor fixed effects. Standard errors are clustered by date, stock, investor, with t-statistics reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5 and 10% levels, respectively.

Panel A. Technical Analysis (TA)

	$\text{Usage}_{i,j,t,n}^{TA}$			
	(1)	(2)	(3)	(4)
Technical Investor$_j$	0.103***	0.075***	0.068***	
	[19.74]	[17.50]	[17.39]	
Swing or Day Trader$_j$		0.020***	0.022***	
		[3.39]	[4.17]	
Long-Term Investor$_j$		-0.028***	-0.023***	
		[-6.10]	[-5.63]	
Professional$_j$		0.047***	0.033***	
		[6.45]	[5.28]	
Novice$_j$		-0.022***	-0.014***	
		[-7.46]	[-5.32]	
Technical$_{i,j,t,n}^{TF-IDF}$			0.799***	0.654***
			[40.24]	[39.56]
Fundamental$_{i,j,t,n}^{TF-IDF}$			-0.632***	-0.498***
			[-26.71]	[-28.51]
Log(# Words$_{i,j,t,n}$)			0.049***	0.043***
			[18.93]	[26.78]
Date FE	No	No	Yes	Yes
Stock FE	No	No	Yes	No
Investor FE	No	No	No	No
Stock × Investor FE	No	No	No	Yes
N	21,641,362	21,641,362	21,641,218	20,630,883
R^2	0.022	0.030	0.084	0.287

Panel B. Fundamental Analysis (FA)

	$\text{Usage}^{FA}_{i,j,t,n}$			
	(1)	(2)	(3)	(4)
Technical Investor$_j$	-0.073***	-0.052***	-0.042***	
	[-12.02]	[-11.12]	[-11.11]	
Swing or Day Trader$_j$		-0.026***	-0.015***	
		[-5.16]	[-3.58]	
Long-Term Investor$_j$		0.038***	0.032***	
		[4.80]	[4.78]	
Professional$_j$		0.034***	0.018***	
		[4.06]	[3.05]	
Novice$_j$		-0.031***	-0.016***	
		[-8.80]	[-5.90]	
$\text{Technical}^{TF-IDF}_{i,j,t,n}$			-0.319***	-0.234***
			[-22.32]	[-23.59]
$\text{Fundamental}^{TF-IDF}_{i,j,t,n}$			0.761***	0.567***
			[16.87]	[17.11]
Log(# Words$_{i,j,t,n}$)			0.137***	0.125***
			[47.15]	[58.24]
Date FE	No	No	Yes	Yes
Stock FE	No	No	Yes	No
Investor FE	No	No	No	No
Stock $\times$ Investor FE	No	No	No	Yes
N	21,641,362	21,641,362	21,641,218	20,630,883
R^2	0.009	0.016	0.161	0.325

Panel C. Other Strategy (OS)

	$\text{Usage}^{OS}_{i,j,t,n}$			
	(1)	(2)	(3)	(4)
Technical Investor$_j$	0.032***	0.017***	0.014***	
	[14.43]	[8.53]	[7.73]	
Swing or Day Trader$_j$		0.011***	0.012***	
		[5.01]	[5.27]	
Long-Term Investor$_j$		-0.016***	-0.013***	
		[-6.94]	[-5.91]	
Professional$_j$		0.022***	0.015***	
		[7.44]	[5.71]	
Novice$_j$		-0.015***	-0.011***	
		[-9.78]	[-7.41]	
$\text{Technical}^{TF-IDF}_{i,j,t,n}$			0.220***	0.185***
			[15.53]	[15.10]
$\text{Fundamental}^{TF-IDF}_{i,j,t,n}$			-0.275***	-0.211***
			[-23.04]	[-20.94]
Log(# Words$_{i,j,t,n}$)			0.029***	0.032***
			[23.86]	[47.78]
Date FE	No	No	Yes	Yes
Stock FE	No	No	Yes	No
Investor FE	No	No	No	No
Stock $\times$ Investor FE	No	No	No	Yes
N	21,641,362	21,641,362	21,641,218	20,630,883
R^2	0.003	0.005	0.021	0.145

Table 3: **Retail Investment Strategy Malleability**

This table reports results from panel regressions on the malleability of retail investment strategies. We examine various factors potentially influencing retail strategy usage: (i) public news releases, (ii) investor strategy performance, and (iii) social feedback on strategy adoption. Panels A, B, and C correspond to these three sets of factors respectively. We focus on investors who have posted at least one strategy-related message in the past three months. The regression is estimated at the individual message level, with the following specification:

$$\begin{aligned} \textit{Usage}^{type}_{i,j,t,n} =& \beta_1 \textit{Public News Releases}_{i,t} + \beta_2 \textit{Investor Strategy Performance}^{type}_{j,t-1} + \\ & \beta_3 \textit{Social Feedback}^{type}_{j,t-1} + \beta_4 \boldsymbol{X_{i,j,t,n}} + \boldsymbol{FE} + \epsilon_{i,j,t,n}, \quad type \in \{TA, FA, OS\} \end{aligned}$$

where $\textit{Usage}^{type}_{i,j,t,n}$ is an indicator equal to one if message n, posted by investor j about stock i on day t, is classified as strategy type TA, FA, or OS. $\textit{Public News Releases}_{i,t}$ represents the indicators for earnings news, analyst news (e.g., recommendations, price targets), or other business news (e.g., mergers, credit ratings, labor issues) for stock i on day t. $\textit{Investor Strategy Performance}^{type}_{j,t-1}$ captures investor j's strategy-specific performance, revealed by investor j's sentiments in messages posted over the prior three months, with High (Low) Performance indicating investor j is assigned to the top (bottom) quartile. $\textit{Social Feedback}^{type}_{j,t-1}$ reflects the total number of likes on strategy-specific messages received by investor j in the prior three months, with High (Low) Likes representing the top (bottom) quartile. $X_{i,j,t,n}$ is a vector of control variables varying across panels. Panel A includes controls for investor j's usage on TA, FA, and OS, as measured by strategy-specific message count, over the prior three months. Panel B adds the indicators of public news releases as controls. Panel C further controls for the prior 3-month investor-level strategy-specific performance. All regressions include date, stock, investor fixed effects. Standard errors are clustered by date, stock, investor, with t-statistics reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5 and 10% levels, respectively.

Panel A. Public News Releases			
	$\text{Usage}^{TA}_{i,j,t,n}$	$\text{Usage}^{FA}_{i,j,t,n}$	$\text{Usage}^{OS}_{i,j,t,n}$
	(1)	(2)	(3)
I(Earnings News)$_{i,t}$	-0.025***	0.048***	-0.011***
	[-15.68]	[20.32]	[-13.14]
I(Analyst News)$_{i,t}$	-0.019***	0.014***	-0.005***
	[-19.20]	[8.79]	[-9.08]
I(Business News)$_{i,t}$	-0.023***	0.014***	-0.007***
	[-20.52]	[10.14]	[-12.71]
Past-Strategy Usage	Yes	Yes	Yes
Date FE	Yes	Yes	Yes
Stock FE	Yes	Yes	Yes
Investor FE	Yes	Yes	Yes
N	15,925,616	15,925,616	15,925,616
R^2	0.204	0.209	0.073

Panel B. Investor Strategy Performance			
	$\text{Usage}^{TA}_{i,j,t,n}$	$\text{Usage}^{FA}_{i,j,t,n}$	$\text{Usage}^{OS}_{i,j,t,n}$
	(1)	(2)	(3)
$\text{High Performance}^{TA}_{j,t-1}$	-0.001	-0.000	-0.000
	[-0.63]	[-0.05]	[-0.16]
$\text{Low Performance}^{TA}_{j,t-1}$	-0.007**	0.008***	-0.002
	[-2.51]	[2.59]	[-0.83]
$\text{High Performance}^{FA}_{j,t-1}$	0.003	-0.002	0.001
	[1.63]	[-1.18]	[0.56]
$\text{Low Performance}^{FA}_{j,t-1}$	0.005***	-0.005***	0.003***
	[3.08]	[-2.88]	[2.95]
$\text{High Performance}^{OS}_{j,t-1}$	-0.005**	0.007**	-0.001
	[-2.22]	[2.23]	[-0.67]
$\text{Low Performance}^{OS}_{j,t-1}$	-0.002	0.002	-0.001
	[-1.10]	[0.71]	[-0.52]
Past-Strategy Usage	Yes	Yes	Yes
Public News	Yes	Yes	Yes
Date FE	Yes	Yes	Yes
Stock FE	Yes	Yes	Yes
Investor FE	Yes	Yes	Yes
N	3,542,296	3,542,296	3,542,296
R^2	0.106	0.091	0.039

Panel C. Social Feedback on Strategy Adoption

	$\text{Usage}^{TA}_{i,j,t,n}$	$\text{Usage}^{FA}_{i,j,t,n}$	$\text{Usage}^{OS}_{i,j,t,n}$
	(1)	(2)	(3)
$\text{High Likes}^{TA}_{j,t-1}$	0.013***	-0.002	-0.003*
	[4.62]	[-0.69]	[-1.85]
$\text{Low Likes}^{TA}_{j,t-1}$	-0.013***	0.007***	0.001
	[-6.38]	[3.35]	[0.67]
$\text{High Likes}^{FA}_{j,t-1}$	-0.008***	0.015***	-0.005***
	[-3.24]	[4.59]	[-2.90]
$\text{Low Likes}^{FA}_{j,t-1}$	0.011***	-0.017***	0.008***
	[4.90]	[-8.15]	[4.14]
$\text{High Likes}^{OS}_{j,t-1}$	-0.009***	-0.005**	-0.000
	[-3.33]	[-2.44]	[-0.24]
$\text{Low Likes}^{OS}_{j,t-1}$	0.006***	0.005**	-0.001
	[2.94]	[2.20]	[-0.98]
Past-Strategy Usage	Yes	Yes	Yes
Past-Strategy Performance	Yes	Yes	Yes
Public News	Yes	Yes	Yes
Date FE	Yes	Yes	Yes
Stock FE	Yes	Yes	Yes
Investor FE	Yes	Yes	Yes
N	3,542,296	3,542,296	3,542,296
R^2	0.106	0.091	0.039

Table 4: **Strategy-Specific Sentiment and Next-Day Stock Returns**

This table reports results from predictive regressions of the next-day stock return with retail sentiments by different strategy types. StockTwits messages are classified into four categories by LLMs: Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), and Non-Strategy (NS). The regression is estimated at the stock-day level, with the following specification:

$$\textit{Return}_{i,t+1} = \beta_1 \textit{Sentiment}_{i,t}^{type} + \beta_2 \textit{Attention}_{i,t} + \gamma \boldsymbol{X_{i,t}} + \boldsymbol{\delta_t} + \epsilon_{i,t+1}, \quad type \in \{TA, FA, OS, NS\}$$

where $\textit{Sentiment}_{i,t}^{type}$ denotes StockTwits investor sentiment toward stock i on day t, separately measured for each strategy type (TA, FA, OS, or NS). Following Cookson et al. (2024a), sentiment is defined as the difference between bullish and bearish message counts across all investors, normalized by their sum:

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

$\textit{Attention}_{i,t}$ captures StockTwits investor attention toward stock i on day t, defined as the percentage of total messages posted on day t that reference stock i on that day, $\frac{\#Messages_{i,t}}{\sum_i \#Messages_{i,t}} (\times 100)$. $\boldsymbol{X_{i,t}}$ represents a vector of control variables including the logarithm of market capitalization, the logarithm of book-to-market, asset growth, gross profit-to-assets, the logarithm of analyst coverage, the logarithm of institutional ownership, maximum daily return in the prior month, abnormal turnover, abnormal news article volume, and returns at five daily lags. All regressions include trading day fixed effects. Standard errors are clustered by trading day, with t-statistics reported in brackets. ***, **, and * indicate 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
Date FE	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 5: **Long–Short Strategies by Strategy Sentiments**

This table reports results from performance evaluation on daily long-short (L/S) trading strategies formed on retail sentiments, categorized by strategy types: Technical Analysis (TA), Fundamental Analysis (FA), and Other Strategies (OS). Following the signal-based strategy construction methodology in Jensen et al. (2023), the L/S strategy takes positions across the entire cross-section of stocks with valid sentiment measures, and its return is calculated based on the deviation of each stock's sentiment from the cross-sectional average:

$$r_t^{L/S,type} = \frac{\sum_{j=1}^{N}(S_{j,t-1}^{type} - S_{t-1}^{type})r_{j,t}}{\frac{1}{2}\sum_{j=1}^{N}\left|S_{j,t-1}^{type} - S_{t-1}^{type}\right|}, \quad \text{where } S_{t-1}^{type} = \frac{1}{N}\sum_{j=1}^{N} S_{j,t-1}^{type}, \quad type \in \{TA, FA, OS\}.$$

We employ the approach of Nagel (2005) to mitigate potential confounding effects on the daily return predictability of retail investor sentiments. Specifically, on each day, we first estimate a cross-sectional regression of daily sentiment measures on a set of daily stock characteristics – investor attention, the logarithm of market capitalization, abnormal turnover, and returns at five lags – and then use the residual sentiments to form the L/S trading strategies. The table summarizes the annualized average daily returns and Sharpe ratios for these sentiment-based L/S strategies. Panel A considers daily raw returns, and Panel B focuses on daily DGTW excess returns (Daniel et al., 1997). For intuitive interpretation, TA and OS sentiments are multiplied by -1 when forming the respective portfolios.

Panel A. Raw Return

	Average (annual, %)	t-statistic	SR (annual)
	(1)	(2)	(3)
TA	9.50	2.91	0.86
FA	6.48	2.04	0.58
OS	10.29	2.92	0.83

Panel B. DGTW Excess Return

	Average (annual, %)	t-statistic	SR (annual)
	(1)	(2)	(3)
TA	10.10	3.21	1.00
FA	7.75	2.53	0.75
OS	8.77	2.55	0.75

Table 6: **Predicting Stock Returns at Longer Horizons**

This table extends the analysis of the next-day return predictability in Table 4 to longer forecasting horizons (up to 15 days ahead). Specifically, we estimate the following predictive regressions at the stock-day level:

$$\textit{Return}_{i,h} = \beta_1 \textit{Sentiment}_{i,t}^{type} + \beta_2 \textit{Attention}_{i,t} + \gamma \boldsymbol{X_{i,t}} + \boldsymbol{\delta_t} + \epsilon_{i,t+1}, \quad type \in \{TA, FA, OS, NS\}$$

where h denotes the three forecasting horizons: $t+1$ to $t+5$, $t+6$ to $t+10$, or $t+11$ to $t+15$. Retail investor sentiments toward stock i on day t ($\textit{Sentiment}_{i,t}^{type}$) are revealed by StockTwits messages classified by LLMs into Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), and Non-Strategy (NS). 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.

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(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
Date FE	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

Table 7: **Return Predictability by Investor Sophistication**

This table reports results from panel regressions examining how investor experience affects the return predictability of retail investor sentiment. Daily sentiment measures, categorized by strategy type (TA, FA, OS, and NS), are interacted with $FracProMsg_{i,t}$, a proxy for the presence of professional investors, defined as the fraction of total messages about stock i on day t posted by self-declared professional investors. Standard errors are clustered by trading day, and t-statistics are reported in brackets. ***, **, and * indicate 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)
$\text{Sentiment}^{TA}_{i,t} \times \text{FracProMsg}_{i,t}$	0.107***	0.128***	0.063*
	[3.03]	[3.78]	[1.95]
$\text{Sentiment}^{FA}_{i,t} \times \text{FracProMsg}_{i,t}$	0.027	-0.026	-0.004
	[0.75]	[-0.77]	[-0.11]
$\text{Sentiment}^{OS}_{i,t} \times \text{FracProMsg}_{i,t}$	0.118***	0.084**	0.085**
	[2.90]	[2.15]	[2.24]
$\text{Sentiment}^{NS}_{i,t} \times \text{FracProMsg}_{i,t}$	0.035	0.032	-0.039
	[1.26]	[1.15]	[-1.45]
$\text{Sentiment}^{TA}_{i,t}$	-0.087***	-0.089***	-0.032
	[-4.00]	[-4.37]	[-1.60]
$\text{Sentiment}^{FA}_{i,t}$	0.047***	0.007	0.001
	[2.77]	[0.40]	[0.04]
$\text{Sentiment}^{OS}_{i,t}$	-0.116***	-0.063***	-0.055***
	[-5.13]	[-2.89]	[-2.64]
$\text{Sentiment}^{NS}_{i,t}$	-0.031**	-0.007	0.016
	[-2.31]	[-0.47]	[1.20]
$\text{FracProMsg}_{i,t}$	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
Date FE	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

Table 8: **Strategy Sentiment, Retail Trading, and Informativeness**

This table reports results from panel regressions examining the contemporaneous relation between retail investor strategy-specific sentiments and retail market order flows, as well as the informativeness of sentiment-driven retail order flows about future stock returns. Panel A shows the relationship between retail sentiments categorized by strategy type – Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), and Non-Strategy (NS) – and aggregate retail market order imbalances (*OIB*). To align closely with market orders, we compute sentiment measures only using StockTwits messages posted during regular trading hours (i.e., 9:30-16:00). The regression is estimated at the stock-day level, with the following specification:

$$\boldsymbol{OIB}_{i,t} = \sum_{type} \beta^{type} \boldsymbol{Sentiment}_{i,t}^{type} + \beta_3 \boldsymbol{Attention}_{i,t} + \gamma \boldsymbol{X_{i,t}} + \boldsymbol{\delta_t} + \epsilon_{i,t}, \quad type \in \{TA, FA, OS, NS\},$$

where the dependent variable, $OIB_{i,t}$, is computed from retail market orders classified by the methodologies proposed in Boehmer et al. (2021) (*BJZZ*) and Barber et al. (2022) (*BJHOS*), respectively. Panel B shows the next-day return predictability of sentiment-driven *OIB* by strategy types. Specifically, we first decompose *OIB* into strategy-specific components by regressing *OIB* on the contemporaneous sentiments of all four strategy types simultaneously. These sentiment-driven *OIB* components, $OIB_{i,t}^{type}$, calculated as $\widehat{\beta}^{type} \times \boldsymbol{Sentiment}_{i,t}^{type}$, and the residual component, $OIB_{i,t}^{resid}$, are then used in the return predictive regression as follows,

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

The explanatory and control variables are defined as in Table 4. All specifications include trading day fixed effects. Standard errors are clustered by stock and trading day in Panel A and by trading day in Panel B. Corresponding t-statistics are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5 and 10% levels, respectively.

Panel A. Retail Market Order Imbalance		
	$\text{OIB}_{i,t}^{BJZZ}$ (%)	$\text{OIB}_{i,t}^{BHJOS}$ (%)
	(1)	(2)
$\text{Sentiment}_{i,t}^{TA}$	0.728***	0.975***
	[17.48]	[24.97]
$\text{Sentiment}_{i,t}^{FA}$	0.566***	0.612***
	[15.02]	[16.33]
$\text{Sentiment}_{i,t}^{OS}$	0.719***	0.792***
	[19.06]	[22.89]
$\text{Sentiment}_{i,t}^{NS}$	0.485***	0.634***
	[13.57]	[18.20]
$\text{Attention}_{i,t}$	0.385***	0.481***
	[2.75]	[3.38]
Stock Characteristics	Yes	Yes
Lagged Returns	Yes	Yes
Date FE	Yes	Yes
N	2,974,934	2,974,934
R^2	0.009	0.012

Panel B. Retail Order Informativeness

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

Table 9: **Strategy Sentiment and Intense Buying by Robinhood Users**

This table reports results from panel regressions examining the contemporaneous relation between retail investor herding on Robinhood and retail sentiments revealed by StockTwits messages, categorized by strategy type: Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), and Non-Strategy (NS). The regression is estimated at the stock-day level, with the following specification:

$$\textit{RH_Herd}_{i,t} = \beta_1 \textit{Sentiment}_{i,t}^{type} + \beta_2 \textit{Attention}_{i,t} + \gamma \boldsymbol{X_{i,t}} + \boldsymbol{\delta_t} + \epsilon_{i,t}, \quad type \in \{TA, FA, OS, NS\}$$

where the dependent variable, $\textit{RH_Herd}_{i,t}$, is an indicator equal to one if stock i is among the top ten stocks ranked by daily percentage growth in Robinhood users holding the stock, provided that at least 100 users held stock i 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 4. All regressions include trading day fixed effects. Standard errors are clustered by stock and trading day, and t-statistics are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5 and 10% levels, respectively.

	RH_Herd$_{i,t}$ (%)				
	(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
Date FE	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

Appendix A.

A.1. Comparing Classification Approaches

Our main approach is to first generate message-level strategy labels using GPT-4 Turbo and then use these labels to fine-tune a BERT model (namely, TechBERT) for our classification task. We compare this method against two alternative approaches: fine-tuning a BERT model on users' self-declared investment approaches and a traditional dictionary-based method.

First, we evaluate the model trained on self-declared labels. Since some StockTwits users self-declare their investment approach, we can fine-tune a BERT model on messages from these users. To do this, we randomly sample 20,000 messages from this group and assign a message-level label based on the user's self-declaration. After fine-tuning a BERT model on these labels, we apply it to classify all StockTwits messages. Figure A.1 presents the predicted probability (investment approach intensity) from this model. The results show that this approach fails to reliably classify messages. For example, the top-left panel shows the distribution of technical analysis (TA) intensity for messages from self-declared technical investors. The predicted probability of TA intensity is relatively low, with most messages scoring below 50%. This poor performance is likely because investors use diverse strategies and often post messages that do not align with their primary, self-declared approach.

Next, we compare our TechBERT model against the self-declared model using a manually-labeled dataset. Figure A.2 displays the TA intensity scores for 500 human-classified messages. The top panels show that TechBERT provides a clear classification: most human-labeled technical messages receive a high TA intensity score, while most non-technical messages receive a low score, creating a distinct separation. In contrast, the bottom panels show that the BERT model trained on self-declared approaches fails to clearly delineate between the two message types. While the average TA intensity is higher for the human-labeled technical messages, the distribution is not bimodal, making it difficult to set a clear classification threshold.

Finally, we assess the dictionary-based, bag-of-words (BoW) approach. Figures A.3 and A.4 present the TA and fundamental analysis (FA) intensity scores calculated using the word lists from Cookson and Niessner (2020). Similar to the self-declared model, the BoW approach does not produce a bimodal distribution. For instance, messages from self-declared technical investors largely receive a BoW TA intensity score close to 0. Likewise, messages from most fundamental investors receive a BoW fundamental score near 0.

Taken together, our analyses highlight the power of using large language models to generate high-quality training data and demonstrate the ability of smaller, fine-tuned models to effectively learn from this data for specific classification tasks.

A.2. Additional Tables and Figures

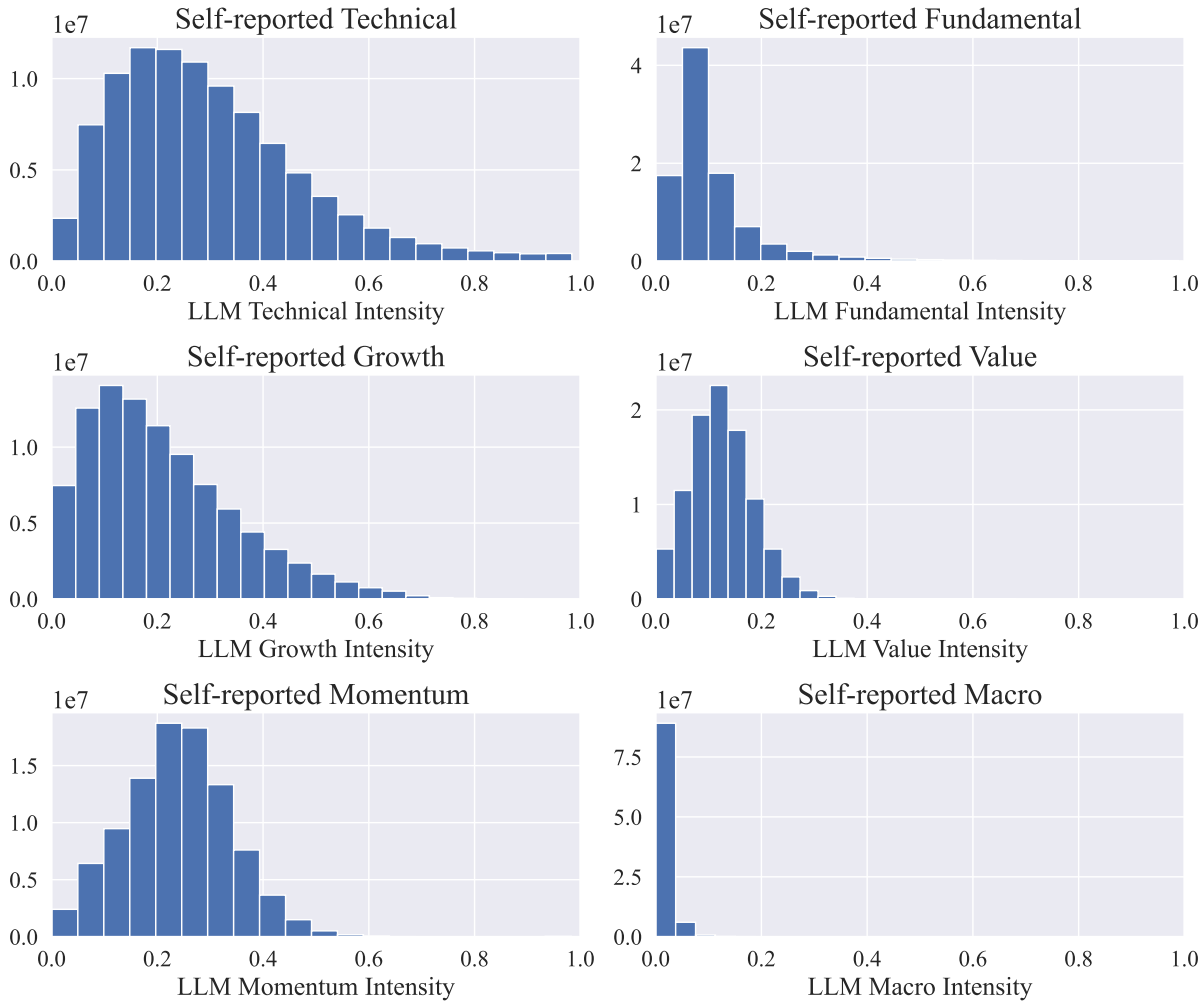


Fig. A.1. **Strategy Classification Using Self-Declared Style + BERT**

This figure presents histograms of message-level classification probabilities from a BERT model fine-tuned on a random sample of messages combined with users' self-declared investment styles. The model assigns each message a probability of belonging to one of six strategy types. For each self-declared strategy type, a panel shows the distribution of BERT-classified probabilities that messages are identified as relying on that same strategy.

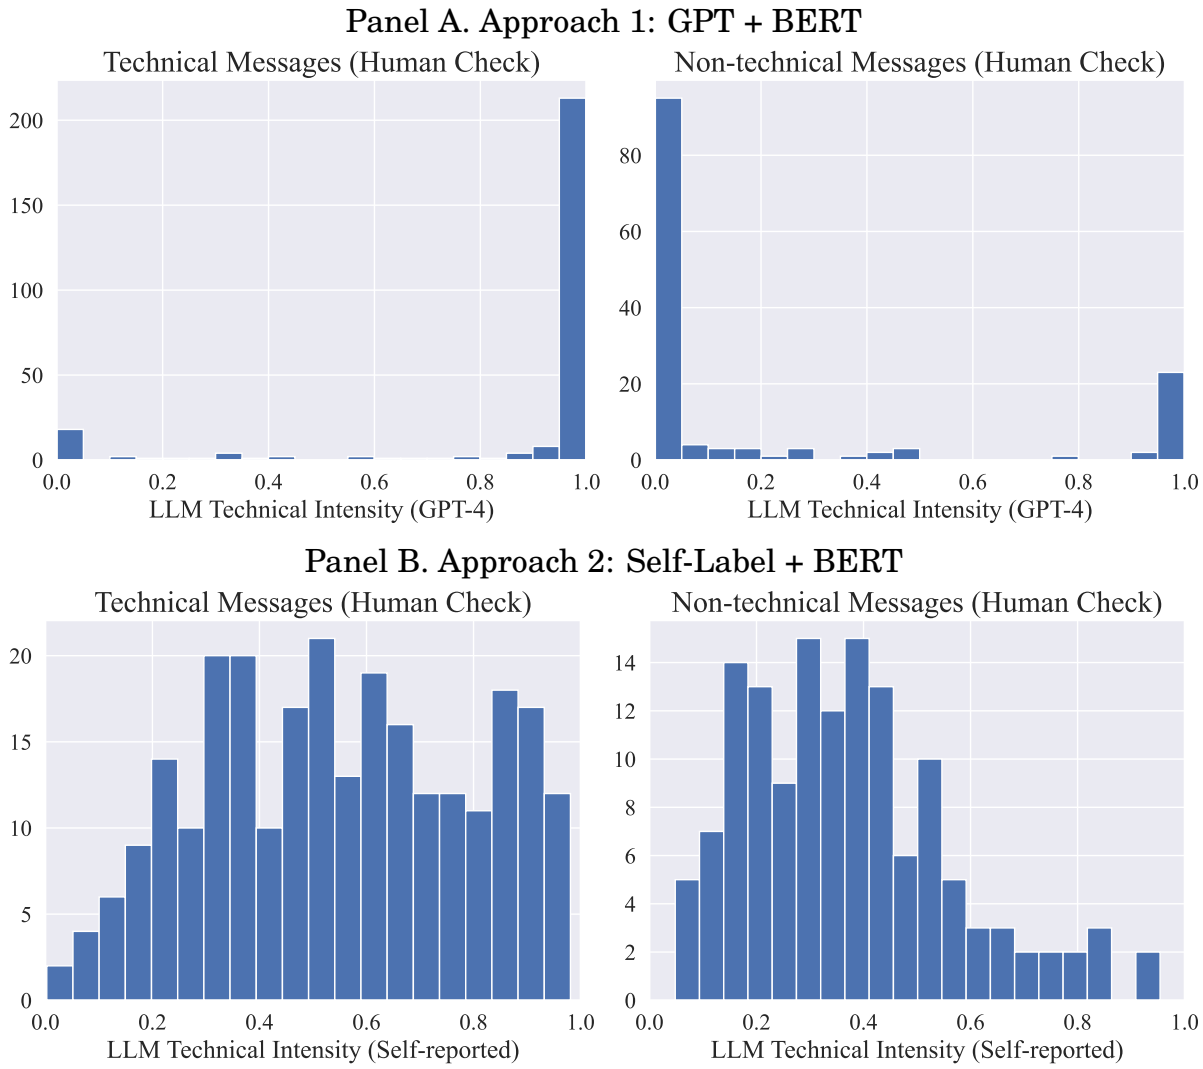


Fig. A.2. **Evaluating Strategy-Classification Methods against Human-Labeled Benchmark**

This figure compares the performance of two alternative strategy-classification methods. The first approach uses GPT-4 Turbo to classify a random sample of messages, and the resulting labels are used to fine-tune a BERT model. The second approach fine-tunes a BERT model using a random sample of messages combined with users' self-declared investment styles. To evaluate performance, we construct a "ground truth" benchmark consisting of two message samples classified by research assistants: (I). Human-Labeled technical messages and (II). Human-labeled non-technical messages. Each method is then applied to classify the benchmark messages and produces a probability that a message relies on technical analysis. Panel A presents histograms of technical-analysis probabilities from the first approach for Sample I (left) and Sample II (right). Panel B presents the corresponding histograms from the second approach.

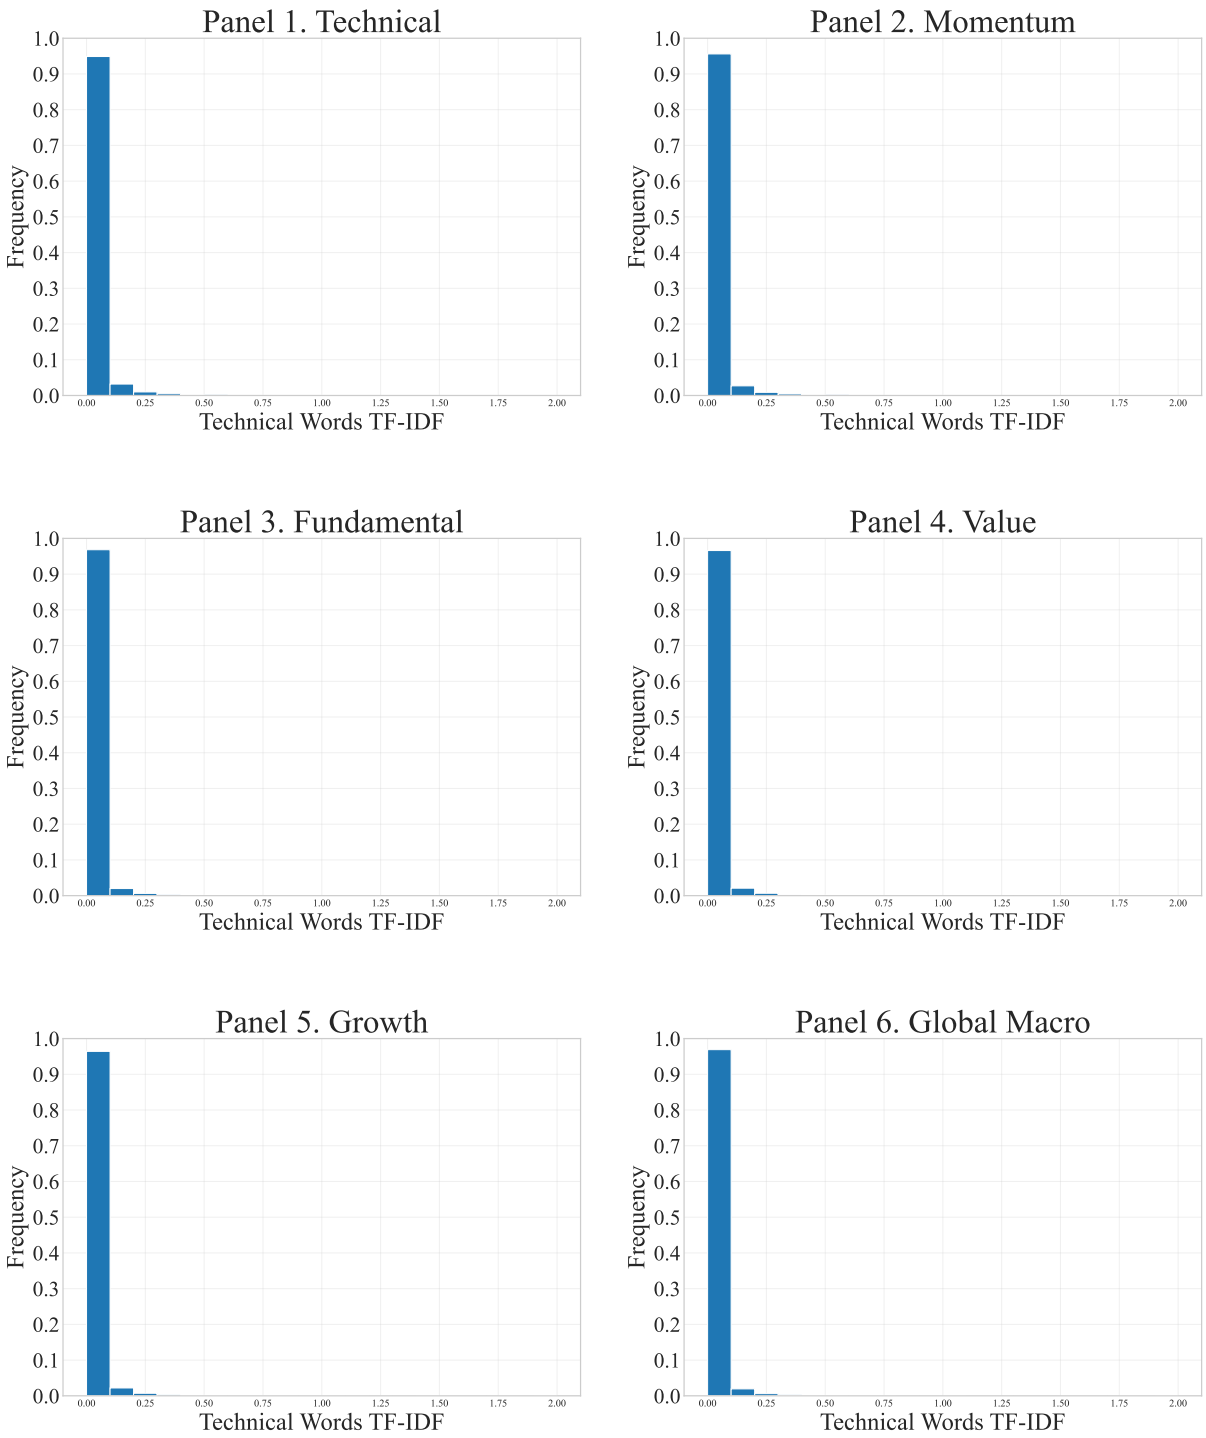


Fig. A.3. **TF-IDF Scores of TA Keywords Across Self-Declared Investment Approaches**

This figure shows technical words TF-IDF scores at the message level. The technical word dictionary is obtained from Cookson and Niessner (2020). Each panel presents the distribution of TA keywords TF-IDF scores for a self-declared investment approach.

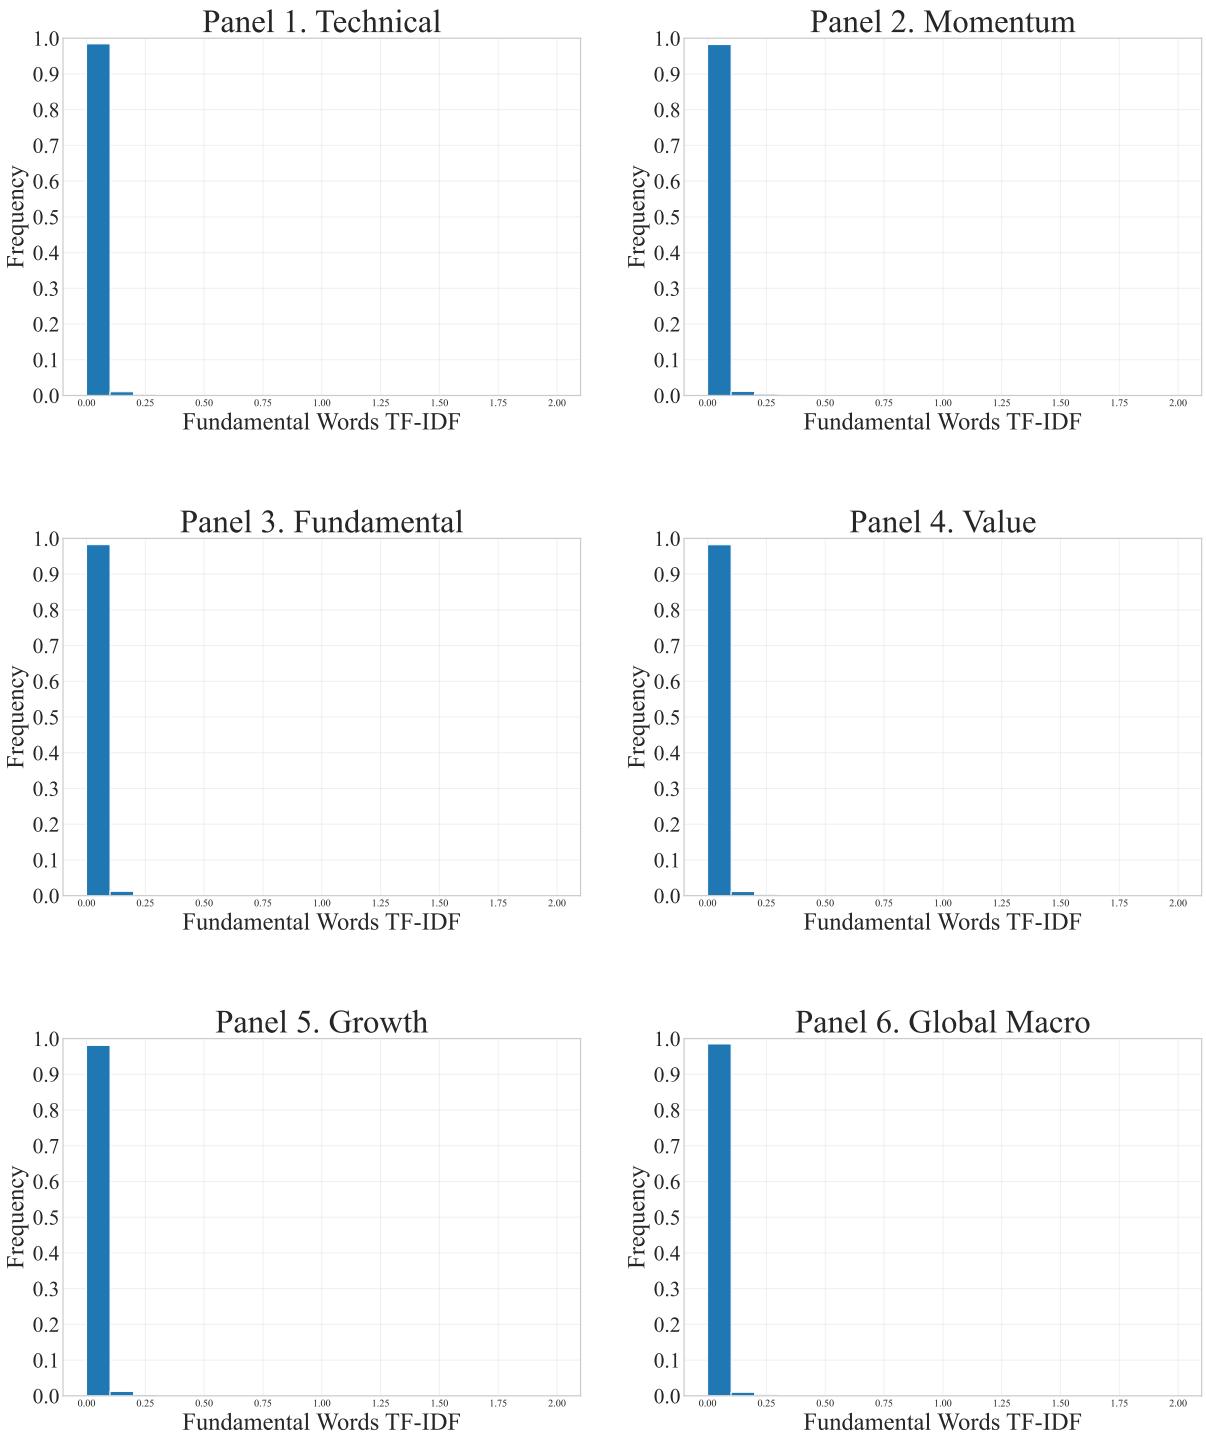


Fig. A.4. TF-IDF Scores of FA Keywords Across Self-Declared Investment Approaches

This figure shows fundamental words TF-IDF scores at the message level. The fundamental word dictionary is obtained from Cookson and Niessner (2020). Each panel presents the distribution of FA keywords TF-IDF scores for a self-declared investment approach.

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

Response to the TA Prompt				
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	-

Response to FA Prompt				
No.	Message	Ticker	Score	Indicator
1	Actually nervous to see $AAPL earnings. People expect too much and realistic is never good enough. Still fundamentally one of best stocks.	AAPL	2	earnings
2	ioDrive2 qualifications a negating effect for $FIO revenues next quarter? BS imo. What about ioDrive which probably takes 2-3 qrtrs??	FIO	2	revenues
3	$CHK company should just put itself up for sale....assets are worth way more than the current stock price....no doubt	CHK	2	assets
4	$NTAP not liking that discussion of non-organic rev was down 9% last yr in 1q	NTAP	2	revenues
5	Piper Jaffray details 10 Apple strengths for share price run up to $1000 - report $AAPL	AAPL	2	analyst-ratings
6	$CSTR People waking up to fact $CSTR has 2 dying businesses–DVDs and Coincashing(anyone ever hear of debit cards and streaming)	CSTR	2	products-services
7	Rising selling margin on full-price goods is a good sign only if folks are buying more of them. Sadly not the case for Penney $JCP	JCP	2	revenues
8	Or that it's trading at 0.5 P/B (historically trades at 1.2-2.0 P/B)? @Thinkb4trading $AIG- does anyone realize the PE ratio is "3"?	AIG	2	assets
9	$TSLA A lot of batteries will be needed in Florida. Quasi Republican Elon, to the rescue, selling batteries to Florida in need. Heard Generac is ready for high demand for batteries.	TSLA	1	products-services
10	$AMC Good news they just finished Filming Honey I Shrunk the Kids 2 !!	AMC	1	products-services
11	$TLRY my first very small position with 850 shares (bought last week) isn't printing yet... time to buy more... this is imo one of the best plays for eoy... enough catalysts in front of us... double digits and more... everybody buying options/shares here?	TLRY	1	products-services
12	Earnings whisper says $SOFI will beat. I'm bullish on the name for growth	SOFI	1	earnings
13	in one week, may see that again, gets us to 335 level everyone is talking about $AAPL	AAPL	0	
14	if $AAPL dips below 435 tmrow, I'm going to jump in with some of the wkly calls - even if they are expensive - and write some more puts too	AAPL	0	

Response to the Strategy Prompt				
No.	Message	Ticker	Score	Category
1	$AAPL looks like it is being pegged to the $600 quarterly option strike	AAPL	1	event-driven strategy
2	$AAPL BTO Jan 2013 $650 Call @ $64; BTC May 2012 $750 Call @ $3.25 & May 2012 $775 @ $1.9 for approx. 20% overnight gains.	AAPL	1	event-driven strategy
3	Watching $AAPL expiring 620 puts, tide could turn fast : $1.25 x 1.39:	AAPL	1	event-driven strategy
4	I think this is a pretty big negative for this stock, could test those 80 cents-$1 lows in 2009. Really is bad news for capital plans $GNW	GNW	1	event-driven strategy
5	Wow, any bond funds that bought the $HGSI convert straight up in Nov must be feeling very good. Now fetching $126 after being as low as $84.	HGSI	1	event-driven strategy
6	$CHK got filled on June Put Spread, got out of weekly put from Friday.. in @ .09 out today @ .23. Now long June $17 and short June $12	CHK	1	event-driven strategy
7	RT @GOODGREED: $AAPL $610 tomorrow as shorts panic to cover...	AAPL	1	event-driven strategy
8	$DNKN- congrats to macro investors here- been pounding the table on this one the last few month- $33 close would be good	DNKN	1	macro
9	“@tunwang: $META huge earnings on mobile, way to go. will get back to above $50.?” Where it should belong higher. Bullish.	META	0	
10	$META if you dumped below $50.... good. burn with the rest pussbags	META	0	
11	$TTWO HUGE block trades: 131355 shares traded - $17.91 @ 3pm yesterday & 45900 shares traded -$17.91 @ 07:50:08 today.	TTWO	0	
12	$T What’s with the recent rise of ATT? It’s gone from around the $34 range to $36+? Someone fill me in please	T	0	
13	$HK WTI Crude down to $94.98 & Brent down to $106.69. Could be impacking Halcon Resouces Corp.	HK	0	
14	$TSLA fortunately my trigger was number hit, their system went down before .. save me $5000 bucks. I would have stopped out regardless	TSLA	0	

Table A.2: **Variable Definitions**

Variable	**Definition**
$\text{Usage}^{type}_{i,j,t,n}$	Indicator variables 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.
$\text{Technical Investor}_j$	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".
$\text{Technical}^{TF-IDF}_{j,i,t,n}$	The TF-IDF weighted score for technical keywords in message n.
$\text{Fundamental}^{TF-IDF}_{j,i,t,n}$	The TF-IDF weighted score for fundamental keywords in message n.
$\text{I(News Coverage)}_{i,t}$	Firm-level daily indicators such as earnings news, analyst news (e.g., recommendations and price targets), and other business news (e.g., mergers, credit ratings, labor issues).
$\text{High (Low) Performance}^{type}_{j,t-1}$	Measures the performance of a hypothetical trading strategy that follows an investor's strategy-specific sentiment over the prior three months, calculated from the subsequent 5-day stock returns after their messages are posted. From this measure, we define High (Low) Performance as an indicator variable for investors whose performance falls in the top (bottom) quartile (i.e., above the 75th or below the 25th percentile) of the cross-sectional distribution.
$\text{High (Low) Likes}^{type}_{j,t-1}$	Reflects the social feedback to an investor for each strategy. We measure an investor's social feedback based on the total number of likes on strategy-specific messages posted over the prior 3 months. High (Low) Likes indicator is defined as the groups of investors based on the 75th and 25th percentiles of their received likes.
$\text{Number of Words}_{i,j,t,n}$	Total number of words in a given message.
$\text{Sentiment}_{i,t}$	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, $\frac{N^{Bullish}_{i,t}-N^{Bearish}_{i,t}}{N^{Bullish}_{i,t}+N^{Bearish}_{i,t}}$ following Cookson and Niessner (2020).
$\text{Sentiment}^{TA}_{i,t}$	Sentiment calculated using messages related to technical analysis.
$\text{Sentiment}^{FA}_{i,t}$	Sentiment calculated using messages related to fundamental analysis.
$\text{Sentiment}^{OS}_{i,t}$	Sentiment calculated using messages that are not related to other strategies.
$\text{Sentiment}^{NS}_{i,t}$	Sentiment calculated using messages that are not related to technical analysis, fundamental analysis, or other strategies.
$\text{Attention}_{i,t}$	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., 2024a).
$\text{OIB}_{i,t}$	Retail marketable volume imbalance on stock i in week t following Boehmer et al. (2021) (BJZZ) or Barber et al. (2023a) (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, calculated as price times shares outstanding.
$\text{Book-to-Market}_{i,t}$	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.
$\text{Analyst Coverage}_{i,t}$	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 strategy-specific 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 strategy-specific sentiments on DGTW excess returns.

	$\text{DGTW}_{i,t+1\to t+5}$ (%)	$\text{DGTW}_{i,t+6\to t+10}$ (%)	$\text{DGTW}_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
$\text{Sentiment}_{i,t}^{TA}$	-0.059***	-0.059***	-0.024*
	[-4.17]	[-4.37]	[-1.81]
$\text{Sentiment}_{i,t}^{FA}$	0.048***	0.000	-0.002
	[3.97]	[0.04]	[-0.18]
$\text{Sentiment}_{i,t}^{OS}$	-0.063***	-0.031**	-0.038***
	[-4.10]	[-2.16]	[-2.70]
$\text{Sentiment}_{i,t}^{NS}$	-0.025**	0.002	0.005
	[-2.43]	[0.19]	[0.54]
$\text{Attention}_{i,t}$	-0.097***	-0.047***	-0.023**
	[-6.14]	[-4.19]	[-2.45]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
Date FE	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 requiring stock-days to have least 10 Stocktwits messages as in Cookson et al. (2024a).

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(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
Date FE	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 4 using the alternative sentiment measures computed based on the maximum entropy method proposed in Cookson and Niessner (2020).

	Return$_{i,t+1}$ (%)				
	(1)	(2)	(3)	(4)	(5)
Sentiment$^{TA}_{i,t}$	-0.024***				-0.022***
	[-3.13]				[-3.18]
Sentiment$^{FA}_{i,t}$		0.010			0.016***
		[1.60]			[2.69]
Sentiment$^{OS}_{i,t}$			-0.032***		-0.031***
			[-3.91]		[-4.03]
Sentiment$^{NS}_{i,t}$				-0.006	-0.002
				[-0.99]	[-0.40]
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
Date FE	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 by Sub-periods**

This table presents the baseline results on the return predictability with retail investors' strategy-specific sentiments by subperiods.

	$\text{Return}_{i,t+1}$ (%)		
	(1)	(2)	(3)
$\text{Sentiment}_{i,t}^{TA}$	0.008	-0.018*	-0.027**
	[0.96]	[-1.72]	[-1.97]
$\text{Sentiment}_{i,t}^{FA}$	-0.005	0.028***	0.022**
	[-0.52]	[3.21]	[2.01]
$\text{Sentiment}_{i,t}^{OS}$	-0.005	-0.026**	-0.033***
	[-0.57]	[-2.49]	[-2.64]
$\text{Sentiment}_{i,t}^{NS}$	-0.001	-0.003	-0.002
	[-0.20]	[-0.30]	[-0.18]
$\text{Attention}_{i,t}$	-0.014**	-0.160***	-0.121**
	[-2.12]	[-4.92]	[-2.36]
Stock Characteristics	Yes	Yes	Yes
Lagged Returns	Yes	Yes	Yes
Date FE	Yes	Yes	Yes
Periods	2010-2015	2016-2019	2020-2023
N	721,296	930,682	1,322,326
R^2	0.093	0.051	0.104

Table A.8: **Retail Order Informativeness and Strategy Sentiment, Longer Horizons**

This table examines the longer horizon return predictability of retail market order imbalance (*OIB*) driven by strategy-specific sentiments as in Table 8.

Panel A. Decomposing BJZZ Retail Order Flow

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(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
Date FE	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

Panel B. Decomposing BHJOS Retail Order Flow

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(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
Date FE	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

Table A.9: **Retail Order Informativeness and Strategy Sentiment, Alternative Measures**

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.

Panel A. Decomposing BJZZ Retail Order Flow

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(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
Date FE	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

Panel B. Decomposing BHJOS Retail Order Flow

	$\text{Return}_{i,t+1\to t+5}$ (%)	$\text{Return}_{i,t+6\to t+10}$ (%)	$\text{Return}_{i,t+11\to t+15}$ (%)
	(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
Date FE	Yes	Yes	Yes
N	2,972,286	2,970,213	2,968,174
R^2	0.105	0.108	0.112

Table A.10: **Standard Deviation of OIB Measures**

This table presents the standard deviation of the OIB measures for TA, FA, OS, NS, and the residual components.

Variable	BJZZ ($\times 100$)	BHJOS ($\times 100$)
OIB^{TA}	0.34	0.47
OIB^{FA}	0.19	0.25
OIB^{OS}	0.25	0.31
OIB^{NS}	0.15	0.26
OIB^{RESID}	27.40	24.68