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

Shuaiyu Chen† Lin Peng ‡ Dexin Zhou§

November 9, 2025

Abstract

We use rich social media data and large language models to classify retail investors' trading strategies, showing that strategy adoption is highly dynamic, and responds to news, past performance, and social feedback. Sentiment from fundamental analysis posts positively predicts future returns, whereas sentiment from technical or other strategy posts negatively predicts returns. Technical sentiment is strongly correlated with net retail buying, especially among Robinhood investors. Retail order flows are substantially more informative when linked to fundamental analysis sentiment. These results demonstrate how retail investors form strategies and the conditions under which strategy adoption enhances or diminishes their order flow informativeness.

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), Lauren Cohen, Zhi Da, Chukwuma Dim, Winston Dou, David Hirshleifer, Byoung-Hyun Hwang, Celine Fei (discussant), Jerry Hoberg, Clifton Green, Lakshmi Naaraayanan (discussant), Marina Niessner, Ron Kaniel, Jinfei Sheng (discussant), John Shim, Chishen Wei (discussant), Dacheng Xiu, Zonghao Yang (discussant), and seminar participants at Baruch College, CUHK, CUHK-Shenzhen, Purdue University, Binghamton University, George Washington University, McMaster University, Rochester Institute of Technology, University of Central Florida, University of Notre Dame, Tsinghua University (PBCSF), 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, RAPS/RCFS European, USC Conference on Social and Behavioral Finance, Wolfe Research AI in Finance Conference, Qube Research & Technologies (QRT) Academy Workshop, and the GSU-MS AI and FinTech Conference.

†chens@darden.virginia.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 emergence of fintech brokerage platforms and social media sites has been accompanied by a dramatic rise in stock market participation by retail investors.[1] This trend highlights the growing importance of retail investors in financial markets, underscoring the need to better understand their behavior. The existing literature presents a multifaceted picture: on one hand, retail investors are often characterized as noise traders subject to behavioral biases (e.g., Barber and Odean, 2000; Barber et al., 2022); on the other, the collective net buying of retail investors has been found to predict higher future stock returns, consistent with informed trading or liquidity provision (e.g., Kaniel et al., 2008; Kelley and Tetlock, 2013; Boehmer et al., 2021).

This divergence in the literature may stem from the heterogeneity in the strategies used by retail investors, a notion supported by Cookson and Niessner (2020), who find that different investment approaches help explain the dispersion in investor beliefs. However, much of the existing work implicitly treats strategy choice as a relatively static investor characteristic. We extend this line of inquiry by positing that the investment approach of an investor is not a fixed trait but a dynamic choice that is responsive to feedback such as past performance and social interactions. Specifically, this paper investigates the factors associated with these strategy shifts and their relationship to performance and market-level outcomes, thereby providing a more nuanced understanding of retail investor behavior.

Studying these dynamic choices requires a data source that can capture the thought formation process of market participants in real time. Social media messages provide a unique window into this process. We use a rich dataset of approximately 100 million messages posted by 840,000 investors on StockTwits, cover-

[1] By 2021, retail trading in the U.S. accounted for almost as much volume as mutual funds and hedge funds combined. Furthermore, at the start of 2023, daily inflows from U.S. retail investors reached a record high of $1.5 billion—more than double the pre-2019 figure. Sources: 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.

ing roughly 7,800 stocks from January 2010 to June 2023. As one of the largest investment-focused social media platforms, StockTwits has become an important data source for recent research on the attention, beliefs, and sentiments of retail investors (e.g., Cookson and Niessner, 2020; Cookson et al., 2023, 2024b).

Extracting investors' thoughts from social media messages can present a significant challenge, as these messages are often colloquial and fragmented, rendering traditional text analysis methods inadequate. To overcome this challenge, we develop and validate a cost-effective, two-stage large language model (LLM) approach. First, we use the GPT-4 Turbo model to generate initial strategy classifications for a small sample of messages. These classifications then serve as training data for a tailored BERT model that we use to classify the full sample. We validate this method against the classifications of human raters and find that it achieves high performance in extracting financial concepts from social media texts, exceeding that of traditional bag-of-words techniques. This process classifies each message into one of four categories: technical analysis (TA), fundamental analysis (FA), other strategies (OS, such as options trading), or no strategy (NS).

Applying this methodology to our dataset reveals that 31% of all messages explicitly reference an investment strategy. Among these strategy-related messages, 44% reference FA, 28% reference TA, and 28% reference OS.[2] These initial findings provide the foundation for our main analysis of the factors related to investors' decisions to adopt or abandon these strategies over time.

Before examining these factors, we first establish that strategy use is, in fact, a dynamic choice rather than a static trait. We test this by assessing how much of the variation in strategy use can be explained by static investor characteristics or a combination of investor and stock fixed effects. While we find that strategy adoption is linked to certain attributes (e.g., TA is more common among self-declared momentum traders, FA among long-term investors), these fixed factors explain only a

[2] Messages may be classified into more than one category, but fewer than 1% mention multiple strategies. The remaining 69% are classified as no specific strategy (NS) and often include memes, catchphrases like "to the moon," or other content unrelated to specific strategies.

limited share of the variation. Even with the inclusion of time and investor×stock fixed effects, a regression model yields an R-squared value of only about 30%. The substantial unexplained variation suggests that strategy use is highly variable. This variability is evident even at the individual level. For example, Figure 1 shows posts from a single self-declared technical strategist, *ACInvestorBlog*, who in one instance discusses earnings and valuations (FA) and in another relies on support levels and volume patterns (TA).

Given this strong evidence of dynamic choice, we investigate the factors that are associated with this evolution. Our analysis uncovers three primary patterns. First, we observe a relationship with information availability: 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 sample 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 is available.

Second, we find a correlation with performance feedback: poor returns from a strategy induce switching, while strong returns from alternative strategies lure investors away. This pattern is consistent with the behavior modeled in Barberis et al. (2018), where investors weigh both price and valuation signals and exogenously shift their emphasis between the signals over time. Our evidence suggests that such switching is not exogenous but is instead endogenous and responsive to the investor's own performance, pointing to opportunities for future models to incorporate this dynamic explicitly.

Third, strategy adoption is related to social feedback. An investor is more likely to continue discussing a strategy if their posts receive positive feedback, measured as the number of "likes". Conversely, the investor shows a tendency to shift to alternative strategies if those strategies receive more likes. These findings support the view in Hirshleifer (2020) that social transmission biases favor folk models that are more heavily cued in the environment. Moreover, the association with social feedback can sometimes be stronger than that with past performance, underscoring the important role of social dynamics in investor behavior. For example, an investor's decision to abandon a TA strategy is 85% more sensitive to receiving a low number of "likes" than to the strategy's poor past performance.

Having examined the factors associated with strategy selection, we next analyze the relationship between these strategies and market-level outcomes. Our first test explores whether different strategies are linked to distinct performance patterns. We calculate stock-level sentiments based on messages from each investment strategy and find sharp heterogeneity in their return predictability. Stocks with bullish TA or OS sentiment earn significantly lower next-day returns, whereas stocks with bullish FA sentiment tend to be followed by positive returns. The economic magnitudes are sizable: a long-short portfolio following FA sentiment yields a 6.48% annualized abnormal return, while contrarian portfolios betting against TA and OS sentiment yield annualized returns of about 9.5% and 10.3%, respectively.[3]

This suggests that, on average, sentiment derived from posts discussing fundamental analysis is informative, as it positively predicts subsequent returns. In contrast, sentiment from technical and other strategies is negatively associated with future returns, a pattern consistent with these sentiments being indicative of behavioral biases. Additionally, the negative return predictability of TA and OS sentiment is less pronounced when a higher fraction of the discussions comes from self-reported professional users. This finding is consistent with the view that investor sophistication may help mitigate these behavioral patterns.

Next, we study whether and how strategy-specific sentiments on social media are related to actual trading activities. Since retail investors are known to be active participants on investor-focused social media, these platforms may offer a valuable real-time reflection of their views. We therefore examine the relationship between strategy sentiments and retail order flows. We first demonstrate that sentiment expressed on StockTwits is closely related to retail investors' trading activity, mea-

[3]These estimates do not incorporate transaction costs, and the high turnover implied by such strategies would likely erode realizable returns. Our objective is not to propose implementable trading strategies but to illustrate the relationship between retail investor discussions on social media and the associated performance heterogeneity.

sured with their net order imbalance. This suggests that discussions on StockTwits provide a real-time window into the decision-making processes of retail investors.

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 is it more often "noise" reflecting behavioral biases? Our approach provides new insights into this debate by allowing for heterogeneity and flexibility in retail investors' strategies. Decomposing retail market order imbalance into components attributable to each social media-based strategy, we find that the predictive ability of retail order imbalance for future returns is strengthened by the FA-related component, but weakened by the TA- and OS-related components.

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 fraction of this information might have been publicly articulated in social media rather than concealed in trading activity. Our evidence also suggests that, despite this information being publicly available, it is not incorporated into prices immediately; instead, prices appear to adjust to it only gradually. The extent to which retail investors disclose information online raises questions about the traditional view that informed traders tend to keep their strategies private. Such disclosures may not always be motivated by pecuniary incentives but could instead reflect social status seeking or other nonpecuniary considerations.

We further focus on retail investors on the Robinhood platform, 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 a strong connection between crowding on technical signals discussed on social media platforms and these speculative surges. The pattern is consistent with

Stein (2009): When traders cannot observe how many others follow the same signals, coordination failures can push prices away from fundamentals. Among retail investors, this phenomenon may be 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 associated with FA sentiments incorporate valuable information and are linked to patterns consistent with price discovery, whereas crowding on technical signals is associated with patterns of inefficiency and subsequent reversals.

Our study contributes to several strands of literature. First, it helps to reconcile the divergent findings on the informativeness of retail trading. While prior work documents that aggregate retail order flow can appear both "smart" and "noisy,"[4] the underlying source of this divergence remains an open question. We advance this literature by demonstrating that the investment strategies investors discuss are a key, underexplored dimension.[5] Specifically, we decompose retail order flow into strategy-specific components and show they have opposing predictive power: the component associated with FA positively predicts future returns, whereas the components linked to TA and OS are negative predictors. This finding is new to the literature and directly addresses this puzzle by linking the "smart money" component of retail flow to fundamental analysis discussions on social media and the "noisy" component to crowding on TA signals discussed by users on the same platforms.

Second, we advance the literature on investment strategy by providing the first large-scale empirical evidence that strategy selection is a dynamic, endogenous

[4]For the former, see, for example, Kaniel et al. (2008), Kelley and Tetlock (2013), Boehmer et al. (2021), and Welch (2022). For the latter studies highlighting behavioral biases or poor performance among retail investors, see Barber and Odean (2000), Kumar and Lee (2006), Barber and Odean (2008), Barber et al. (2022), Bryzgalova et al. (2023), and Barber et al. (2023b).

[5]Existing research has explored heterogeneity along dimensions such as experience (Seru et al., 2010) and gender (Barber and Odean, 2001).

process.[6] Prior work has either treated strategy as a fixed investor characteristic (Cookson and Niessner, 2020) or assumed that shifts between strategies are exogenous (e.g., Barberis et al., 2018). Our key contribution here is to document that strategy switching is systematically associated with three distinct feedback mechanisms: past performance, the availability of firm-specific news, and, notably, social feedback in the form of "likes." By demonstrating this endogenous response, we provide an empirical microfoundation for theoretical models of investor behavior and open new avenues for research into the drivers of strategy adoption.

Finally, our study makes both a methodological and a substantive contribution to the literature on social finance and asset pricing.[7] Methodologically, we pioneer a scalable, two-stage LLM framework to accurately extract structured investment strategies from millions of unstructured social media posts, a task previously considered challenging.[8] Building upon prior work that leverages social media to study the investor thought process (e.g., Chen et al., 2025a), our method moves beyond traditional textual analysis to provide more flexibility and accuracy. Substantively, this new methodology allows us to uncover novel results. Complementing prior evidence on confirmation bias in signal reception (e.g., Cookson et al., 2023), we provide evidence consistent with social transmission bias (Hirshleifer, 2020) by showing that investors' strategy choices are significantly related to peer approval, an effect sometimes stronger than that of past returns.

Furthermore, our strategy-level sentiment data reveal new market dynamics, showing how discussions of speculative strategies like TA are strongly associated with Robinhood herding events and subsequent price reversals. This approach

[6]Our work builds on a rich literature studying fundamental analysis (e.g., Porta et al. (1997); Abarbanell and Bushee (1997)) and technical analysis (e.g., Brown and Jennings (1989); Jegadeesh (1991); Brock et al. (1992); Jegadeesh and Titman (1993); Blume et al. (1994); Lo et al. (2000); George and Hwang (2004); and more recently, Han et al. (2013); Han et al. (2016); Jiang et al. (2023); Murray et al. (2024).

[7]Prior studies establish that social networks shape investor beliefs and affect stock returns and trading volume. 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), and the review by Cookson et al. (2024b).

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

opens new avenues for understanding when retail participation enhances price discovery versus when it may be linked to market instability, a heterogeneity with direct implications for platform regulation and investor protection.[9]

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

[9] As an example of its utility, our methodology empowered a related study (Chen et al., 2025b) to show that AI-driven profits can be generated by trading against retail TA sentiment, a result aligned with the theoretical work of Dou et al. (2024).

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

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 StockTwits data with the CRSP stock universe, we apply several filters 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 retain only messages explicitly referencing exactly one ticker symbol. We then exclude all messages from any user who posted more than 1,000 messages on a single day at any point, 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 subsection, we introduce our methodology for extracting two key textual variables from StockTwits messages: the underlying investment rationale (i.e., trading strategy) and the directional sentiment (i.e., bullishness/bearishness). We first describe our use of LLMs to identify the trading strategies referenced in individual messages. We then explain the procedure for obtaining sentiment measures for each message. Finally, we describe the aggregation process used to construct the stock-level, strategy-specific sentiment measures.

2.2.1. *Identifying Trading Strategies from Messages*

StockTwits messages, like much social media content, are short, filled with abbreviations and colloquial expressions, and often contain non-standard variations in spelling. 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 knowledge distillation approach (Hinton, 2015) enables us to leverage the high accuracy of a large "teacher" model to train a smaller, more efficient "student" model.[11]

This methodology offers several advantages. First, although our validation tests find that GPT-4 Turbo classifications closely align with human judgments, applying GPT-4 Turbo to the full dataset is prohibitively expensive. In contrast, a fine-tuned BERT model provides a far more efficient alternative.[12] Moreover, the two-step procedure outperforms natural alternatives such as a purely dictionary-based approach or an alternative two-step design that trains BERT using user-declared investment approaches. Appendix Section A.1 provides further details of these comparisons.

Our procedure classifies each message into one or more strategy categories: technical analysis (TA), fundamental analysis (FA), or other strategies (OS). Messages not assigned any strategy tag are categorized as no strategy (NS). We employ three independent binary classifiers: a TA-classifier, an FA-classifier, and a general strategy-classifier. A message is tagged as TA or FA if its probability score from the respective classifier exceeds 95%.[13] A message is tagged as OS if it is identified by

[11]See Gu et al. (2023) for a recent review of knowledge distillation 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, which has 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]For example, we estimate that classifying all messages in our dataset with GPT-4 Turbo would cost in excess of $500,000.

[13]TA and FA classifications are not mutually exclusive because the classifiers operate independently. About 0.4% of messages are classified as both TA and FA. We use them in calculating measures for both the TA and FA.

the general strategy-classifier (with a probability > 95%) but is not tagged as either TA or FA. A message is classified as no strategy (NS) only when it fails to meet the criteria for TA, FA, or OS.

Technical Analysis Strategy Classification To illustrate our classification procedure, we begin by identifying messages that refer to technical analysis (TA). We repeat this procedure to identify other strategies. 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 those with 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.[14]

We then use GPT's outputs to fine-tune a BERT model (henceforth TA-BERT) to predict whether a message uses TA.[15] Cross-validation indicates that TA-BERT 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

[14] 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 instructed a PhD research assistant to conduct manual classification. The research assistant was instructed to classify the messages based on their knowledge of technical analysis strategies. The assistant's classifications align with GPT's classifications 90% of the time.

[15] The BERT model is finetuned as a classifier. We treat GPT responses 1 and 2 as positive examples (coded as 1) and 0 as negative examples (coded as 0) in the fine-tuning process.

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 self-declared investment horizons, respectively. A key property of TA-BERT's predictions is that most messages fall in either the very low-probability region ($< 5\%$) or very high-probability 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 apply the same two-step procedure to identify messages containing fundamental analysis (FA) by revising the GPT-4 Turbo prompt. We then fine-tune a specialized FA-BERT model to identify FA-related messages. We use the following prompt:[16]

> You have a deep understanding of the language of social media and fi-

[16]The fundamental analysis topics are based on the topic classifications of financial news articles provided by Ravenpack.

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

Similar to the validation of TA classification, we visualize FA-BERT's classifications in Figure 5 and Figure 6, grouped by self-declared investment approaches and self-declared investment horizons, respectively. As with TA-BERT, we observe that FA-BERT's classification is aligned with the self-declared investment approaches. For example, Figure 5 shows that investors who self-report using fundamental, value, and growth approaches 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.

Finally, to identify messages containing any investment strategy, we use the following prompt:

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

As discussed earlier, we classify a message as containing an "other strategy" (OS) if it refers to an investment strategy but not to either TA or FA.

Overall, about 31% of all StockTwits messages include explicit references to trading strategies. Of these strategy-related messages, FA is the most common, appearing in 44% of them, while TA and OS each appear in 28%.

2.2.2. *Message Sentiment*

We first measure sentiment using the user-declared sentiment flag. 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.[17]

To validate the model, we apply the classifier to messages with sentiment flags 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 confirm our main findings remain both quantitatively and qualitatively unchanged when using the maximum entropy approach of Cookson and Niessner (2020) to impute missing sentiment (see Appendix Table A.6).

We then aggregate sentiments at the firm-strategy level. Following Cookson et al. (2024a), we define sentiment 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

[17] 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). We avoid GPT for sentiment to mitigate look-ahead risk. Instead, BERT's pre-training corpus (books/Wikipedia) does not include forward-looking financial text and is therefore unlikely to introduce look-ahead bias (Google BERT team, 2018).

Panel C reports pairwise correlations between sentiment measures.[18] 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 sentiments derived from different investment approaches capture distinct information.

2.3. *Other Variables*

Our analysis also incorporates 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 day, following 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 at least 100 users held the stock at the end of the prior week (Barber et al., 2022).

Our key control variables include 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, 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. Finally, we consider firm characteristics. These include *Market Capitalization*, *Book-to-Market*, *Asset Growth*, and *Gross Profit-to-Asset*. The information environment and institutional ownership are proxied by *Analyst Coverage*, representing the number of IBES equity analysts, and *Institutional Ownership*, calculated as the fraction of shares outstanding 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

[18]For stock-days without any messages, sentiment is set to 0.

summary statistics in the firm-day sample. Notably, the median firm capitalization is 1.2 billion, which is comparable with that reported in Giannini et al. (2018). This number indicates that our sample tilts toward larger firms, suggesting our results are not limited to small firms with limited liquidity.

3. Retail Investment 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 investment approaches, horizons, and message attributes, much variation remains, even after controlling for stock$\times$investor fixed effects. Building on this finding, we then turn to the potential drivers of this flexibility in retail strategy usage, focusing on the roles of public news releases, strategy performance, and social feedback on strategy adoption.

3.1. *Investor and Stock Attributes*

As shown in Figures 2, 3, 5, and 6, our LLM-based strategy classification broadly corresponds to investors' self-declared investment approaches and horizons, but it also reveals substantial malleability. For instance, many investors who self-identify as technical traders often post messages referencing fundamental signals, and vice versa. This malleability 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 messages from users with available biographical information.[19] Our regression is specified as follows:

$$\mathit{Usage}^{type}_{i,j,t,n} = \beta_1 \boldsymbol{X}^{investor}_{j} + \beta_2 \boldsymbol{Z}^{message}_{i,j,t,n} + \boldsymbol{FE} + \epsilon_{i,j,t,n}, \quad type \in \{TA, FA, OS\}, \tag{1}$$

where $Usage^{type}_{i,j,t,n}$ is an indicator variable set to one if message n, posted by investor

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

j about stock i on day t, is classified by our fine-tuned BERT models as referring to one of the strategy types: 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 (i.e., *Technical Investor* equal to one if an investor's self-reported approach is either *Technical* or *Momentum*), investment horizon (i.e., *Long-Term Investor*), and experience level (i.e., *Novice* or *Professional*).

$\boldsymbol{Z}_{i,j,t,n}^{message}$ represents a vector of message-specific attributes, including message length and the frequency of TA or FA-related terms based on the dictionary of TA and FA 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.

The regression specifications also include an extensive set of fixed effects ($\boldsymbol{FE}$) at different levels: date, investor, stock, investor$\times$stock. These fixed effects account for unobserved heterogeneity and quantify the share of variation 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 variable of interest is investors' self-reported investment approach. As expected, self-reported technical investors (i.e., *Technical* or *Momentum*) are 10.3 percentage points more likely (a 73.5% increase compared to the sample mean) to mention TA in their messages relative to those self-reported as using other approaches (i.e., *Fundamental*, *Value*, or *Growth*). However, the R-squared is only 2.2%, indicating that static self-reported profiles have limited power in explaining the substantial variations in strategy choices.

Column (2) adds investor-level indicators for investment horizons and experience. Short-term investors (i.e., *Day/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 discuss TA than the benchmark group (i.e., *Intermediate*), whereas *Novices* are significantly less likely to do so.

Column (3) controls for message length and traditional BoW technical and fundamental scores, as well as date and stock fixed effects to account for unobserved heterogeneity in TA adoption across stocks and time. We find that the 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 broadly aligned with the BoW approach. However, in the comparison of strategy classification methodologies presented in Appendix Section A.1, the traditional BoW approach appears to perform relatively poorly when classifying social media messages into different strategy types. Importantly, the additional controls, together with the stock and date fixed effects, only moderately raise the R-squared from 3.0% to 8.4%, indicating that stock- and time-specific factors are not the primary drivers of retail TA adoption.

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

We next extend our analysis to FA usage ($Usage^{FA}$) in Panel B. We find that self-reported technical investors are less likely to adopt FA. Long-term experienced investors are more inclined to discuss FA factors in their posts. The inclusion of investor$\times$stock fixed effects increases the R-squared by about 16 percentage points, suggesting that retail investors' strategy choices may vary with static stock characteristics. However, as with TA usage, the R-squared is still only around 32.5%, indicating that investor-specific dynamic factors may play a crucial role in FA usage.

In Panel C, we examine the usage of other strategies ($Usage^{OS}$), reflecting investment approaches unrelated to TA or FA.[20] The patterns for OS exhibit similarities to those of TA: OS usage is more prevalent among self-declared technical, short-term, and experienced retail investors.

Taken together, these findings demonstrate that while the message-level strategy choices are broadly consistent with investors' self-declared approaches, investors adapt their strategies fluidly across time and stocks. This pattern thus underscores the dynamic nature of retail investors' decision-making.

3.2. Strategy Malleability

Having established the considerable variability in investors' strategy choices, we now analyze its potential drivers. We address three key questions. First, how are firm-specific information flows related to investors' strategy usage? For example, when there is greater news coverage of fundamentals, such as earnings releases, analyst recommendations, or other value-relevant disclosures, do investors shift toward fundamental analysis and away from technical analysis 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; Huang, 2019). 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). It is therefore plausible that retail investors adjust their strategy choices in response to their previous experience with the performance of various strategies.

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

[20] These strategies are captured by our fine-tuned BERT model designed to identify general trading strategies, but are not identified as TA or FA by TA- or FA-BERT.

(2020) and Akçay and Hirshleifer (2021), folk models that generate a greater excitement ("buzz") are more likely to spread, leading more agents to adhere to them.[21]

To answer these questions, we extend the earlier analyses to examine various factors that potentially influence 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:

$$
\begin{aligned}
Usage_{i,j,t,n}^{type} = {} & \beta_1 Public\ News\ Releases_{i,t} + \beta_2 Investor\ Strategy\ Performance_{j,t-1}^{all\ types} \\
& + \beta_3 Social\ Feedback_{j,t-1}^{all\ types} + \beta_4 \boldsymbol{X}_{i,j,t,n} + \boldsymbol{FE} + \epsilon_{i,j,t,n}, \quad type \in \{TA, FA, OS\}.
\end{aligned}
\tag{2}
$$

$Public\ News\ Releases_{i,t}$ 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 then calculate the sentiment-weighted average of the 5-day compounded DGTW-adjusted 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. The control vector X includes the frequency of usage in the three strategies over

[21] 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 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 FA usage but negatively associated with TA and 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.[22] In column (1), we analyze the determinants of TA usage. We find that investors who experience poor performance when using TA 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 performance when using FA.

In column (2), we examine FA usage. Similar to TA, FA usage decreases following poor performance when using FA but increases after subpar performance when using TA. 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 in shaping investors' strategy usage. We introduce the total number of likes an investor received for posting messages of a given strategy type as a measure of social feedback, alongside controls

[22]In this analysis, we require valid performance metrics in the prior quarter, leading to a significant loss of observations.

for past strategy usage, performance determinants, and other covariates.

In column (1), where the dependent variable is TA usage, we find that positive social feedback for an investor's TA messages increases the investor's future TA usage by 1.3 percentage points, whereas negative social feedback reduces it by the same magnitude. Furthermore, we observe that social feedback also induces strategy switching: an investor's TA usage rises when receiving negative feedback for posting FA- or OS-related messages and declines when the feedback for those messages is positive.

A similar pattern holds for FA usage: investors increase their reliance on FA strategies following positive feedback on their FA-related posts or negative feedback on their TA or OS posts. The case of OS usage is slightly different. While an investor's propensity to adopt OS strategies responds to the feedback received on TA or FA posts, it is not significantly related to the feedback on OS posts themselves. One possible explanation is that, unlike the more clearly defined styles of TA and FA, OS encompasses a broad range of other strategies, making the measure of social feedback on OS posts noisier and less informative than that for FA or TA.

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 it 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—a magnitude that is 85% larger than the one we document for 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 suggest that strategy choice is not a static trait but rather a dynamic process shaped by public information, personal investment outcome, and social reinforcement. The finding that social feedback appears to be a more powerful driver than personal performance provides novel insight into investor learning on social networks and highlights the importance of the social dimension in retail investment.

4. Performance 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). Yet, how investors' strategic choices translate into performance remains an open question. Our finding that investors flexibly switch between strategies, combined with evidence that investors' self-declared approaches drive the divergence of opinions in asset returns (Cookson and Niessner, 2020), motivates us to further investigate whether strategy choice at the message level systematically affects investment outcomes. In this section, we fill this gap by evaluating the performance of distinct retail investment strategies identified through our LLM-based classification of StockTwits messages.

4.1. Daily Return Predictability

We assess next-day return predictability using strategy-specific sentiment measures. Following Cookson et al. (2024a), we estimate the stock-day predictive regression:

$$Return_{i,t+1} = \beta_1 Sentiment_{i,t}^{type} + \beta_2 Attention_{i,t} + \gamma \boldsymbol{X}_{i,t} + \boldsymbol{\delta}_t + \epsilon_{i,t+1}, \tag{3}$$

where $Return_{i,t+1}$ denotes stock i's return on the next trading day $t+1$. $Sentiment_{i,t}^{type}$ (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). $Attention_{i,t}$ 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 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.[23]

These results provide direct evidence of heterogeneity in the informativeness of retail strategies. The negative return predictability of TA and OS sentiments

[23]Table A.7 in the Appendix presents subsample tests, showing that our findings on return predictability are driven by the more recent period (2016–2023) and are absent in the earlier years (2010–2015). The stronger return predictability in recent years may partly reflect the rise of retail trading and the heightened activity on investment-focused social media platforms, which allow us to better capture these effects.

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 that it depends critically on the type of strategy emphasized.

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.[24] 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.

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

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

.

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—under DGTW performance adjustments (Daniel et al., 1997), as shown in Panel B, where the TA (FA) L/S strategy yields an annualized DGTW-adjusted 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 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.[25]

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 associated with persistent behavioral biases rather than value-relevant information. Investors 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 and OS sentiment is generally on the wrong side of the market, previous studies find certain technical indicators consistently generate abnormal returns (e.g., Jegadeesh and Titman, 1993; Han et al., 2013, 2016; Jiang et al., 2023; Murray et al., 2024). This raises the possibility 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 hypothesis, 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

[25] 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.

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 the "*Wisdom of Crowds*" that fundamental signals are informative even when discussions are dominated by novice traders (see, e.g., Chen et al., 2014; Welch, 2022).

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.

5. Retail Strategies, Trading, and Informativeness

Our analysis thus far suggests that StockTwits investor sentiments associated with different investment strategies yield markedly different investment outcomes. A natural question that arises is whether investors merely "talk the talk" without actually "walking the walk." In other words, do social media discussions translate into actual 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 sentiments are representative of retail investor trading activities by examining their relationship to 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. (2023a), 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:

$$\textit{OIB}_{i,t} = \sum_{type} \beta_1^{type} \times \textit{Sentiment}_{i,t}^{type} + \beta_2 \textit{Attention}_{i,t} + \gamma \boldsymbol{X}_{i,t} + \boldsymbol{\delta}_t + \epsilon_{i,t}. \tag{4}$$

To align closely with market orders, we construct the sentiment measures 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 $\textit{OIB}^{BJZZ}$. A one–standard deviation increase in TA (FA) sentiment corresponds to a 0.39 (0.32) percentage point increase in retail order imbalance. Column (2) confirms these results using $\textit{OIB}^{BHJOS}$, 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.[26]

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, $OIB_{i,t}^{TA}$, $OIB_{i,t}^{FA}$, $OIB_{i,t}^{OS}$, and $OIB_{i,t}^{NS}$, defined as $\widehat{\beta}_1^{type} \times \textit{Sentiment}_{i,t}^{type}$, represent the components of *OIB* attributable to each sentiment type. The residual from this regression, $OIB_{i,t}^{Resid}$, captures the portion of retail order flows uncorrelated with StockTwits sentiments, which may reflect either information withheld from social media or pure noise.

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

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, order imbalances attributable to TA and OS *negatively* predict future stock returns, indicating that these strategies are associated with lower informativeness of retail orders. Second, the order imbalance attributable to FA positively predicts returns, consistent with this component enhancing informativeness. Third, NS-driven order imbalance exhibits 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.

Regarding economic magnitudes, a one-standard-deviation increase in FA-driven OIB leads to a 1.1-bp ($= 0.19 \times 0.0056$) increase in next-day returns. By comparison, a similar increase in the residual component of OIB is associated with a 4.7-bp ($= 0.173 \times 0.27$) increase.[27] 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 23.4% of that withheld) is still economically meaningful. 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 tend to be informed traders or are primarily driven by behavioral biases. By decomposing retail order flows into strategy-specific components, we highlight an important heterogeneity: Retail order flow linked to fundamental signals discussed on social media is associated with positive future returns, a pattern consis-

[27]The standard deviations of FA- and residual-OIB measures are 0.0056 and 0.27, respectively.

tent with price discovery, whereas order flow linked to technical or other speculative signals on social media moves 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,[28] 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. We investigate the extent to which StockTwits sentiment explains Robinhood herding.

Following Barber et al. (2022) and Welch (2022), we employ RobinTrack data from May 2018 to August 2020. As in Barber et al. (2022), we define episodes of intense buying (i.e., 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:

$$RH\_Herd_{i,t} = \sum_{type} \beta_1^{type} Sentiment_{i,t}^{type} + \beta_2 Attention_{i,t} + \gamma \boldsymbol{X_{i,t}} + \boldsymbol{\delta_t} + \eta_i + \epsilon_{i,t}. \quad (5)$$

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 is

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

significantly related to herding, with TA sentiment showing the strongest relationship. The effect is economically meaningful. A one–standard deviation increase in TA sentiment raises the probability of herding by 0.062 percentage points—an 18% increase relative to the unconditional herding frequency of 0.35 percentage points. By comparison, the corresponding effects are 0.038 percentage points (an 11% increase) for FA sentiment and 0.033 percentage points (a 10% increase) 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 is consistent 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 can amplify 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 naive 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 retail investors' strategy choices, sentiment, and trading behavior. By extracting investment strategy choices from retail investors' messages on StockTwits, we show that strategy choices are time-varying and are linked to factors such as public news releases, investors' past performance, and social feedback. Technical Analysis (TA) related posts become more prevalent when firm-specific news is scarce, while Fundamental Analysis (FA) gains traction as news flow intensifies. An investor's usage of a particular strategy responds systematically to the investor's prior performance with the strategy and to the feedback received from peers.

Our empirical analysis also reveals sharp differences in the effectiveness of strategies discussed on StockTwits. Retail TA sentiments on StockTwits *negatively* predict future returns, especially when the discussions are dominated by less sophisticated users, whereas FA sentiments on the same venue *positively* predict future returns. While StockTwits sentiments across all strategy types are positively associated with retail net buying, TA sentiment exhibits the strongest connection to intense buying episodes on Robinhood. Moreover, we find that retail net order flows associated with FA-sentiment are informative about future returns, while those linked to TA and OS sentiment predict returns in the wrong direction.

Overall, our paper demonstrates that combining social media data with LLMs provides a powerful new lens on retail trading. Our analysis reveals how retail strategies, as reflected in their social media discussions, can be linked to outcomes that both inform financial markets and are consistent with patterns of market instability. Looking ahead, our approach offers opportunities for future research, such as investigating how retail behavior interacts with institutional trading in dynamic market environments. It also provides a useful perspective for policy makers by helping them 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, 2001, Boys will be boys: Gender, overconfidence, and common stock investment, *The quarterly journal of economics* 116, 261–292.

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, Hailiang, Byoung-Hyoun Hwang, and Zhuozhen Peng, 2025a, Why do investors like short-leg securities? evidence from a textual analysis of buy recommendations, *The Review of Financial Studies* hhaf068.

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

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* .

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.

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* .

Huang, Xing, 2019, Mark twain's cat: Investment experience, categorical thinking, and stock selection, *Journal of Financial Economics* 131, 404–432.

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, 1993, 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, 2024, Expected Returns and Large Language Models, *Available at SSRN* .

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.

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.

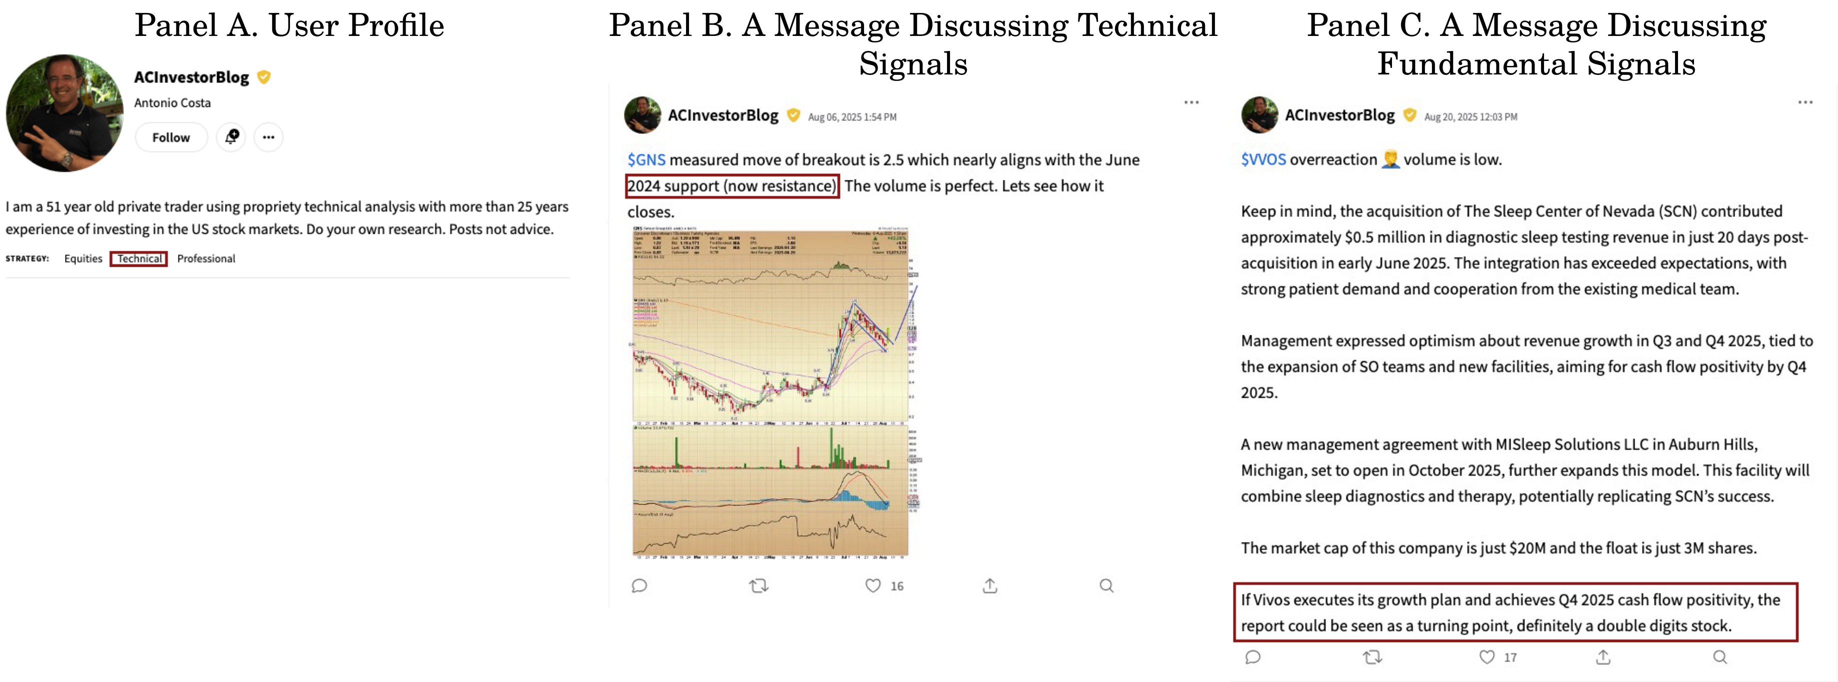

Fig. 1. **Retail Investors' Strategy Malleability: An Illustration**

This figure presents anecdotal evidence on the malleability of retail investors' investment strategies. The screenshots, sourced from StockTwits, feature a user, ACInvestorBlog, who self-identifies as a "*Technical*" strategist (as indicated in the user profile in Panel A). Panel B shows that this user discussed technical signals in a post on 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. In contrast, Panel C shows that the user discussed fundamental indicators in a post on 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.*"

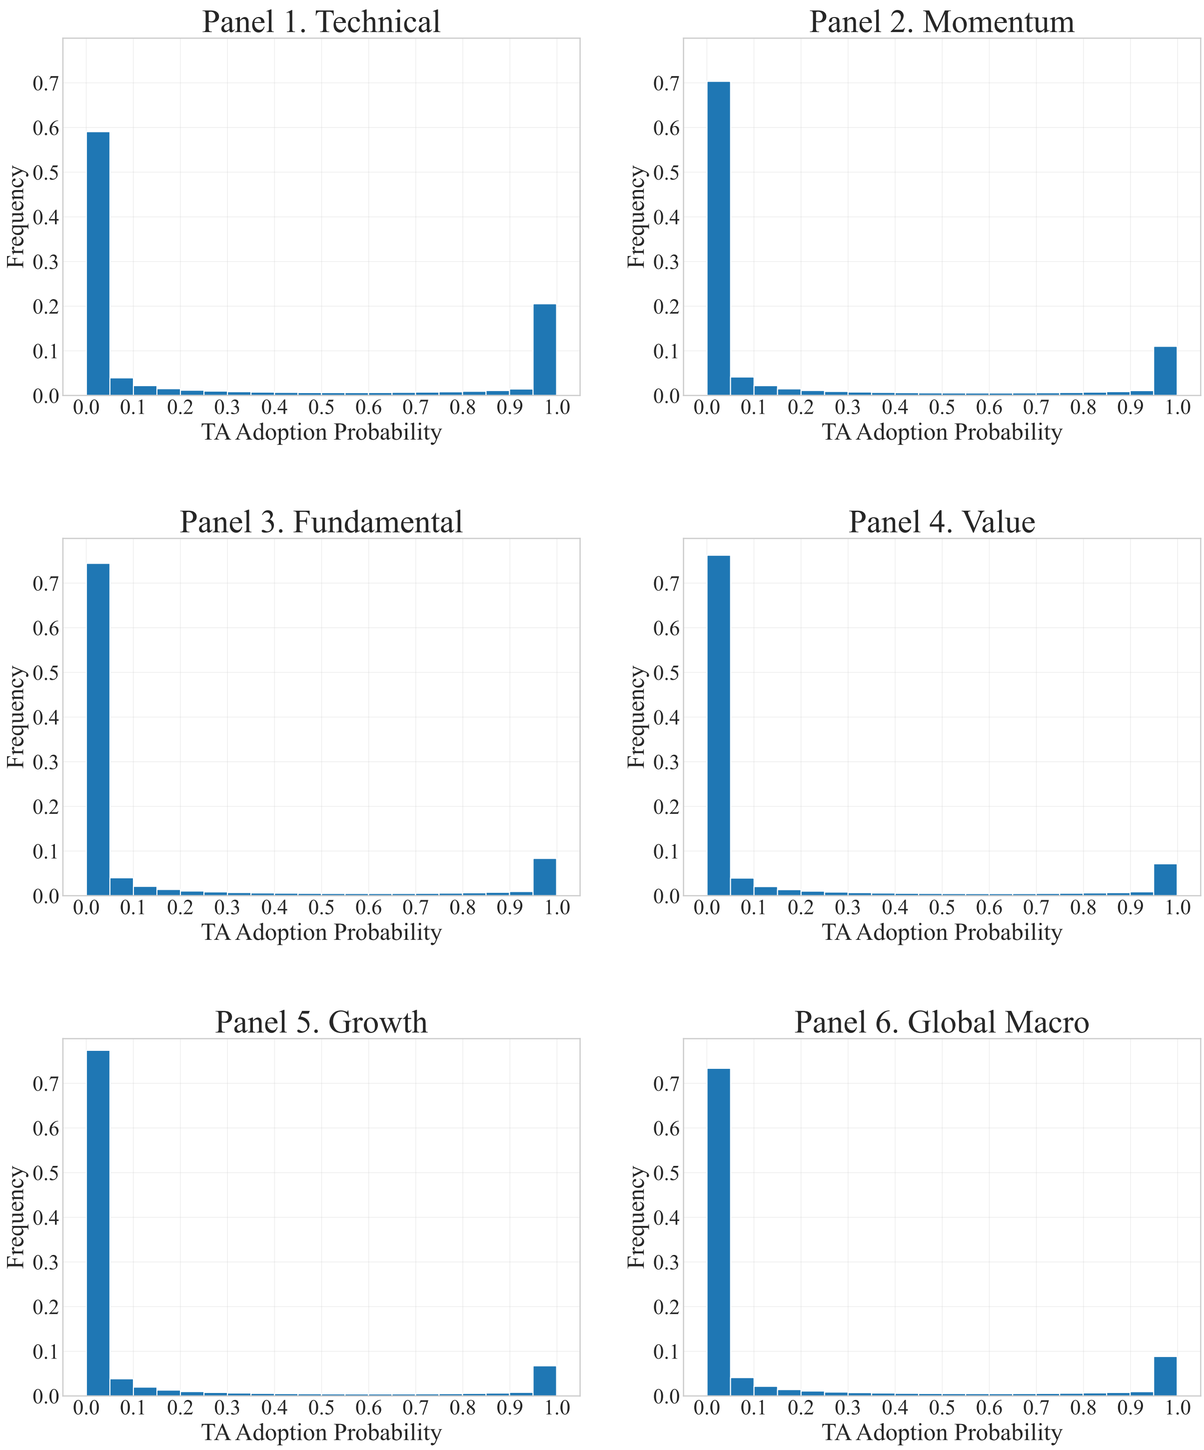

Fig. 2. **Technical Analysis (TA) Adoption Probability Across Self-Declared Investment Approaches**

This figure plots the distribution of *Technical Analysis (TA) Adoption Probability* over StockTwits messages, grouped by users' self-declared investment approaches. 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 approach group, the histogram reports the frequency of messages across levels of *TA Adoption Probability*.

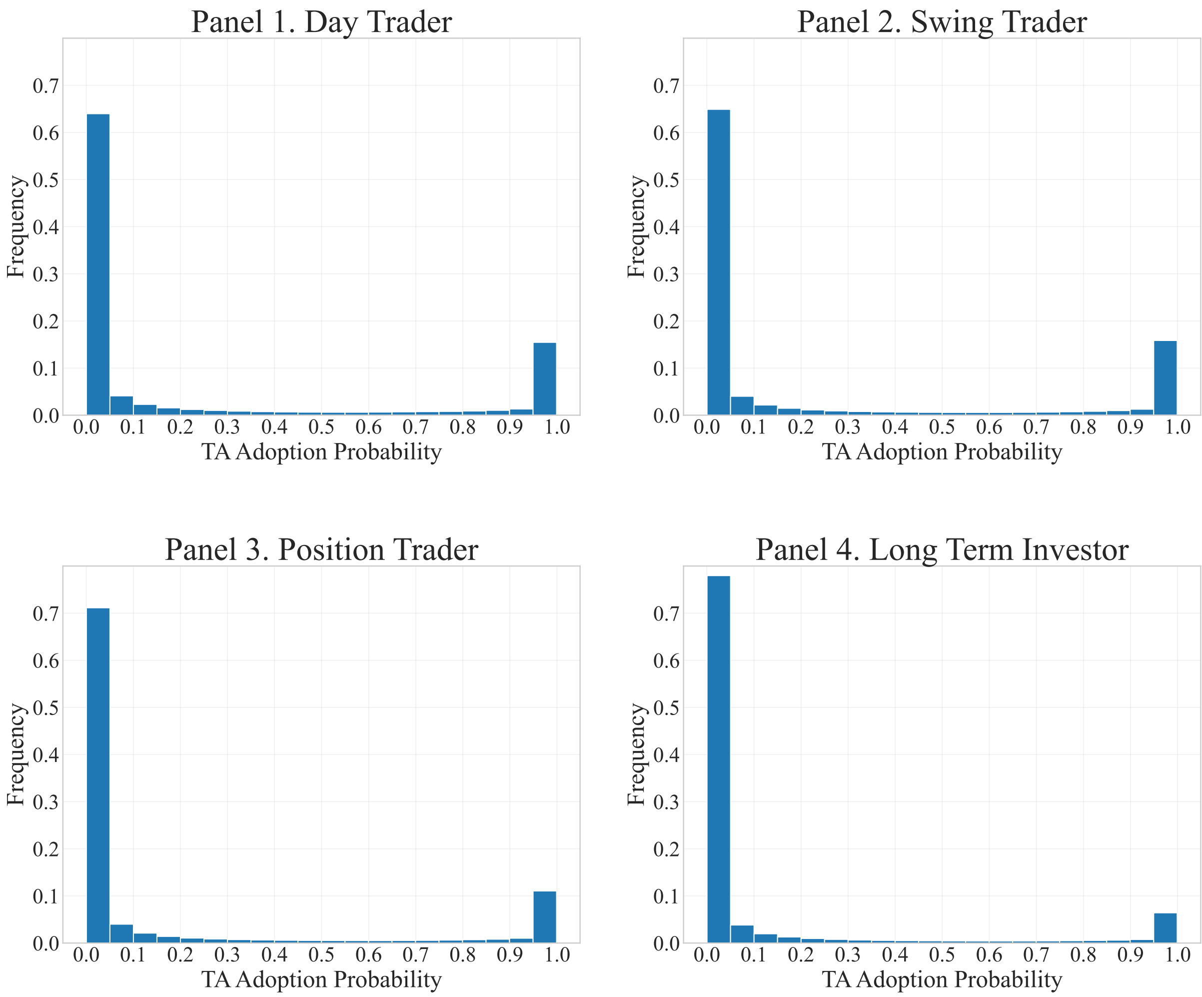

Fig. 3. **Technical Analysis (TA) Adoption Probability Across Self-Declared Investment Horizons**

This figure plots the distribution of *Technical Analysis (TA) Adoption Probability* over Stock-Twits 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 group, the histogram reports the frequency of messages across levels of *TA Adoption Probability*.

Panel A. Word Cloud of Unigrams in Technical Messages

Panel B. Word Cloud of Bigrams in Technical Messages

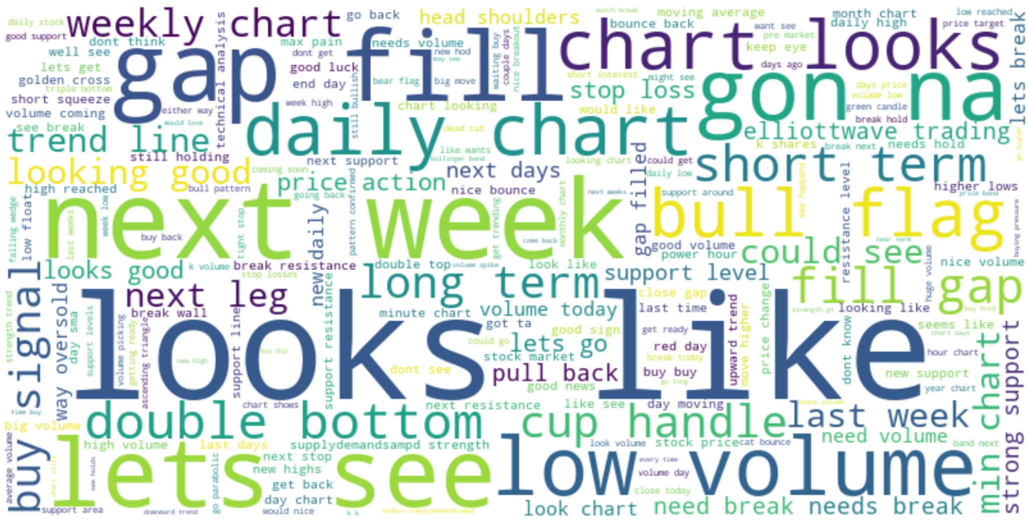

Fig. 4. **Word Clouds of Technical Analysis (TA) Messages**

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.

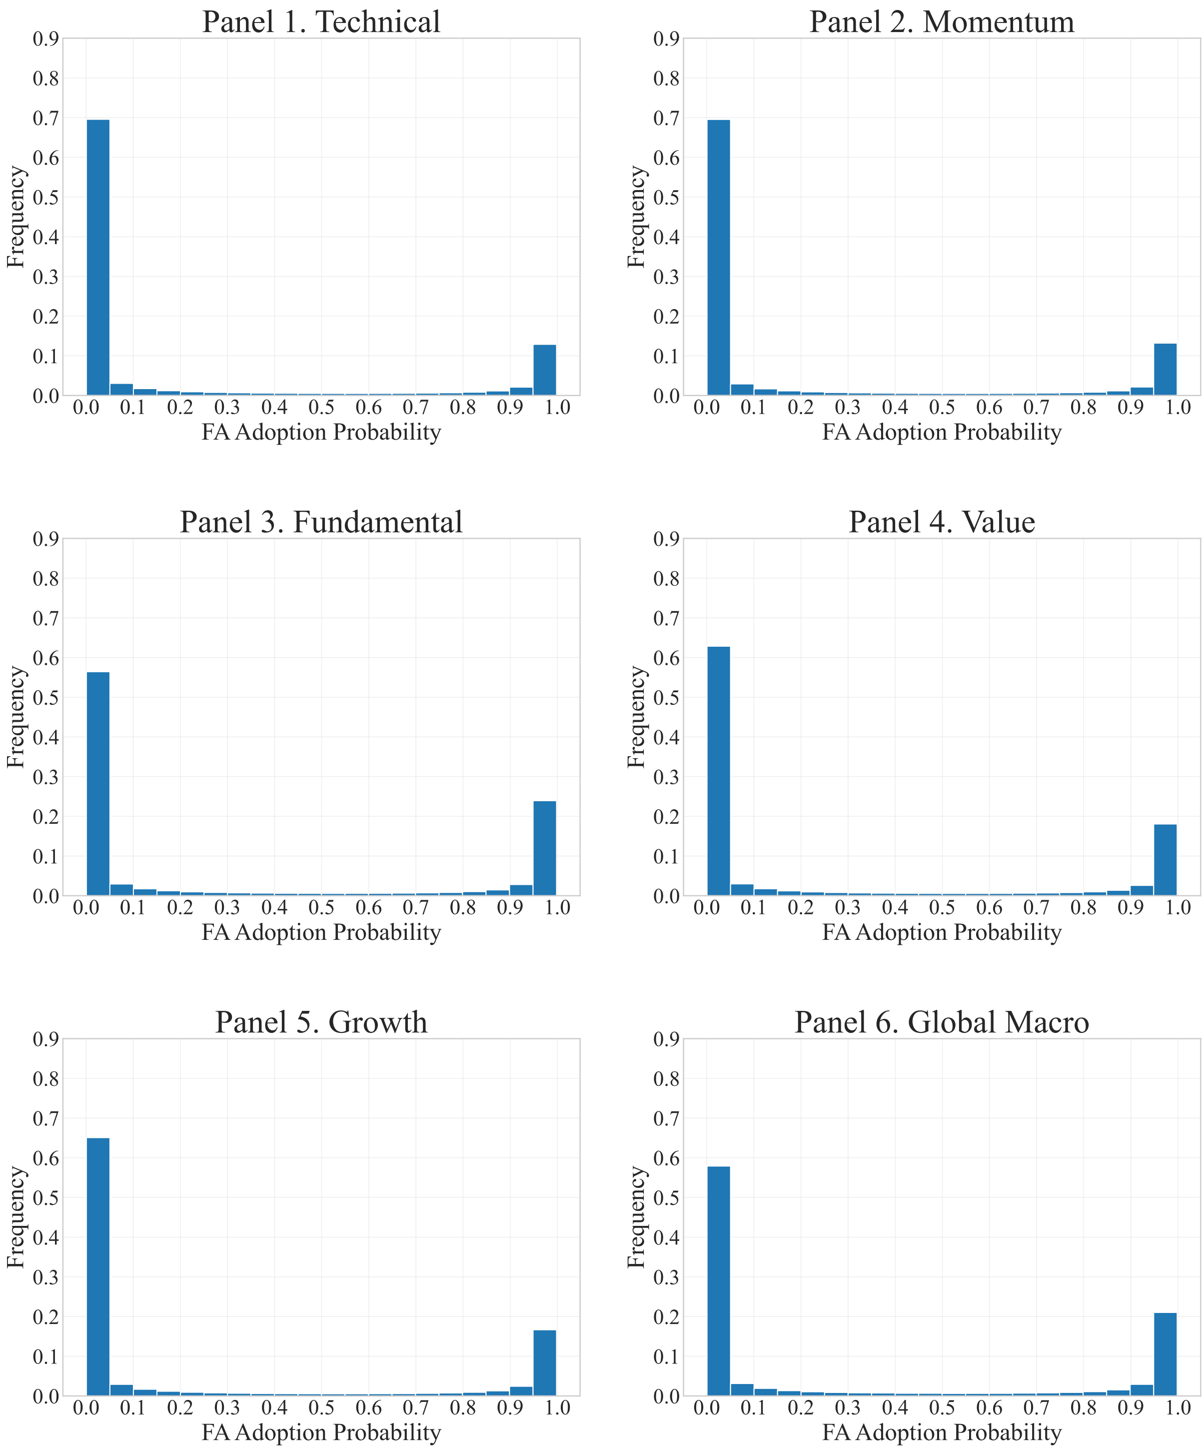

Fig. 5. **Fundamental Analysis (FA) Adoption Probability Across Self-Declared Investment Approaches**

This figure shows the distribution of *Fundamental Analysis (FA) Adoption Probability* over StockTwits messages, grouped by users' self-declared investment approaches. Each message receives a probabilistic score from our fine-tuned *FA-BERT* model, where higher values indicate a greater likelihood that the message employs fundamental analysis. For each investment approach group, the histogram reports the frequency of messages across levels of *FA Adoption Probability*.

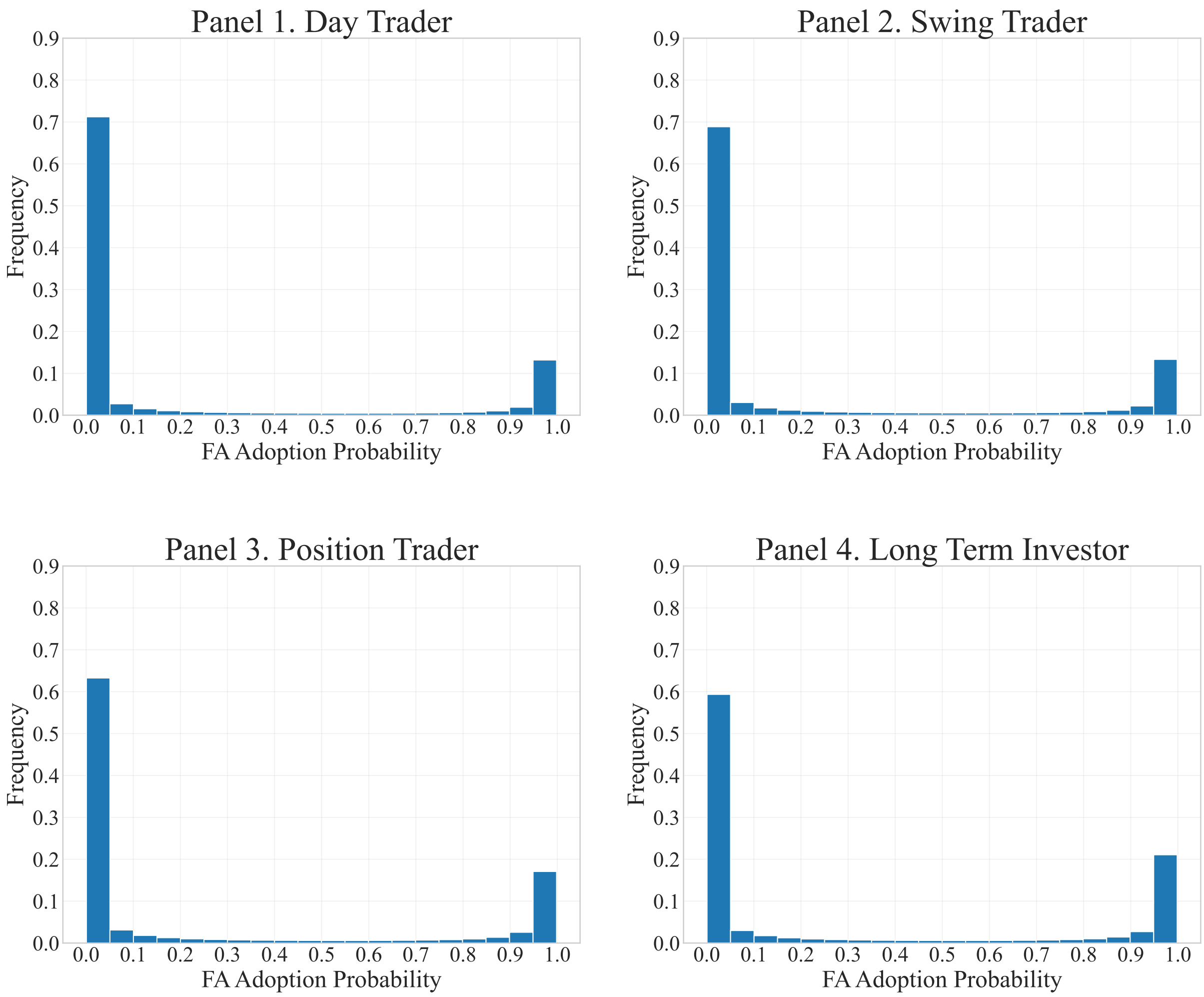

Fig. 6. **Fundamental Analysis (FA) Adoption Probability Across Self-Declared Investment Horizons**

This figure presents the distribution of *Fundamental Analysis (FA) Adoption Probability* over StockTwits messages, 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 fundamental analysis. For each investment horizon group, the histogram reports the frequency of messages across levels of *FA Adoption Probability*.

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 (FA) 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 valid self-reported biographical data. The reported variables include the indicators of LLM-classified investment strategies (Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS)), the indicators of self-declared investor profiles (Technical Investor, Long-Term Investor, Swing/Day Trader, Professional, and Novice), the message length (number of words), and the TF-IDF measures for keywords in technical and fundamental analyses, respectively, based on the word lists from Cookson and Niessner (2020). Panel B presents summary statistics for the stock-day sample, including sentiment measures derived from 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 capitalization, book-to-market, asset growth, gross profit-to-assets, analyst coverage (number of analysts), institutional ownership (IO), the maximum daily return in the prior month (MAX), abnormal turnover, and abnormal news article volume. Panel C reports correlations between sentiment measures 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
RH_Herd								
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
Analyst Coverage	2,974,934	9.74	7.00	8.98	1.00	3.00	15.00	23.00
Institutional Ownership	2,974,934	0.60	0.69	0.32	0.09	0.33	0.86	0.96
MAX	2,974,934	0.08	0.06	0.09	0.02	0.03	0.10	0.16
Abnormal Turnover	2,974,934	-0.11	-0.11	0.63	-0.81	-0.44	0.22	0.61
Abnormal News 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 estimation results from panel regressions relating retail investors' usage of investment strategies, classified by LLMs on StockTwits messages, to the two sets of covariates: (i) self-declared investor attributes, and (ii) message-specific attributes. Panels A, B, and C focus on 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:

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

where $Usage^{type}_{i,j,t,n}$ is an indicator variable equal to one if message n, posted by investor j about stock i on day t, is classified by LLMs into strategy type TA, FA, or OS. Investor-level characteristics ($\boldsymbol{X}^{investor}_j$) include self-reported attributes such as Technical Investor, Long-Term Investor, and Professional Investor. Message-specific attributes ($\boldsymbol{Z}^{message}_{i,j,t,n}$) include the TF-IDF measures 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, and investor, with corresponding t-statistics in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Technical Analysis (TA)

	$\text{Usage}^{TA}_{i,j,t,n}$			
	(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]	
$\text{Technical}^{TF-IDF}_{i,j,t,n}$			0.799***	0.654***
			[40.24]	[39.56]
$\text{Fundamental}^{TF-IDF}_{i,j,t,n}$			-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(# $\text{Words}_{i,j,t,n}$)			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 Strategy Malleability**

This table reports estimation results from panel regressions examining the malleability of retail strategies. We consider the three sets of factors potentially influencing retail strategy usage: (i) public news releases, (ii) investor strategy performance, and (iii) social feedback on strategy. Panels A, B, and C correspond to these three sets, 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:

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

where $Usage^{type}_{i,j,t,n}$ is an indicator equal to one if message n, posted by investor j about stock i on day t, is classified by LLMs into strategy type TA, FA, or OS. $Public\ News\ Releases_{i,t}$ represents the indicators for earnings news, analyst news (e.g., recommendations, price targets), or other business news (e.g., credit ratings, labor issues) for stock i on day t. $Investor\ Strategy\ Performance^{type}_{j,t-1}$ captures investor j's *own* strategy-specific performance, revealed by investor j's sentiments in messages posted in the prior three months, with High (Low) Performance indicating investor j being assigned to the top (bottom) quartile on each day. $Social\ Feedback^{type}_{j,t-1}$ 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 on each day. $X_{i,j,t,n}$ is a vector of control variables varying across panels. Panel A includes controls for investor j's usage on TA, FA, and OS, measured by strategy-specific message count in 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, and investor fixed effects. Standard errors are clustered by date, stock, and investor, with corresponding t-statistics 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 $\text{News}_{i,t}$)	-0.025***	0.048***	-0.011***
	[-15.68]	[20.32]	[-13.14]
I(Analyst $\text{News}_{i,t}$)	-0.019***	0.014***	-0.005***
	[-19.20]	[8.79]	[-9.08]
I(Business $\text{News}_{i,t}$)	-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	$\text{Usage}^{TA}_{i,j,t,n}$	$\text{Usage}^{FA}_{i,j,t,n}$	$\text{Usage}^{OS}_{i,j,t,n}$
	(1)	(2)	(3)
High $\text{Likes}^{TA}_{j,t-1}$	0.013***	-0.002	-0.003*
	[4.62]	[-0.69]	[-1.85]
Low $\text{Likes}^{TA}_{j,t-1}$	-0.013***	0.007***	0.001
	[-6.38]	[3.35]	[0.67]
High $\text{Likes}^{FA}_{j,t-1}$	-0.008***	0.015***	-0.005***
	[-3.24]	[4.59]	[-2.90]
Low $\text{Likes}^{FA}_{j,t-1}$	0.011***	-0.017***	0.008***
	[4.90]	[-8.15]	[4.14]
High $\text{Likes}^{OS}_{j,t-1}$	-0.009***	-0.005**	-0.000
	[-3.33]	[-2.44]	[-0.24]
Low $\text{Likes}^{OS}_{j,t-1}$	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 estimation results from predictive regressions of next-day stock returns with retail investor sentiments by different strategy types. StockTwits messages are classified by LLMs into the four strategy categories: 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 scores toward stock i on day t, separately measured for each strategy type (TA, FA, OS, or NS). Following Cookson et al. (2024a), the sentiment score is defined as the difference between bullish and bearish message counts across all investors, normalized by their sum:

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

$\textit{Attention}_{i,t}$ denotes 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. $\boldsymbol{X}_{i,t}$ 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, the maximum daily return in the prior month, abnormal turnover, abnormal news article volume, and the five daily return lags. 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.

	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 Retail Strategy Sentiments**

This table reports results from performance evaluation on the daily long-short (L/S) trading strategies, separately formed on retail sentiments categorized by three 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 scores, and its return is calculated based on the deviation of each stock's sentiment score from the cross-sectional average:

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

We employ the approach of Nagel (2005) to mitigate potential confounding effects on the daily return predictability of retail sentiments. Specifically, we first estimate a cross-sectional regression on each day of daily sentiment scores on a set of daily stock characteristics – investor attention, the logarithm of market capitalization, abnormal turnover, and five return lags – and then take the residual sentiments to form these L/S strategies. The table summarizes the annualized average daily raw returns or abnormal returns, as well as the Sharpe Ratio (SR) or Information Ratio (IR), for these sentiment-based L/S strategies. Panel A considers raw returns, and Panel B focuses on DGTW-adjusted returns (Daniel et al., 1997). TA and OS sentiment scores are multiplied by -1 in the strategy construction.

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-Adjsuted Return

	Average (annual, %)	t-statistic	IR (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 return predictability analysis in Table 4 to longer predictive horizons (up to 15 days ahead). We estimate the following predictive regressions at the stock-day level:

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

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 are revealed by StockTwits messages classified by LLMs into Technical Analysis (TA), Fundamental Analysis (FA), Other Strategies (OS), and Non-Strategy (NS). The control variables ($\boldsymbol{X}_{i,t}$) are the same as in Table 4. All specifications 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 regressions examining how investor sophistication influences the return predictability of retail strategy sentiments. Daily retail sentiment scores, categorized by strategy type (TA, FA, OS, and NS), are interacted with $FracProMsg_{i,t}$, a proxy for the presence of professional retail investors, defined as the fraction of total messages about stock i on day t posted by self-declared professional StockTwits users. The control variables are the same as in Table 4. All specifications 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)
$\text{Sentiment}^{TA}_{i,t} \times \text{FracProMsg}_{i,t}$	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 Future Returns**

This table reports results that analyze the relationships among retail strategy sentiment, retail order imbalance, and future returns. Panel A examines the link between strategy sentiment and retail order flows. We estimate the following stock-day level regression:

$$\textbf{\textit{OIB}}_{i,t} = \sum_{type} \beta^{type} \times \textit{Sentiment}_{i,t}^{type} + \beta_3 \textit{Attention}_{i,t} + \gamma \boldsymbol{X}_{i,t} + \boldsymbol{\delta}_t + \epsilon_{i,t}, \quad type \in \{TA, FA, OS, NS\}.$$

$OIB_{i,t}$ is computed from retail market orders in TAQ data classified using the methodologies proposed in Boehmer et al. (2021) (*BJZZ*) and Barber et al. (2023a) (*BHJOS*), respectively. *Sentiment*^{type} corresponds to the strategy-specific sentiment variables, *Sentiment*^{TA}, *Sentiment*^{FA}, *Sentiment*^{OS}, and sentiment from non-strategy messages, *Sentiment*^{NS}. To align sentiment measures with retail orders, we compute sentiment scores using only messages posted during regular trading hours (9:30-16:00). Panel B examines the relationship between the retail order flows linked to distinct strategies (TA, FA, OS, and NS) and future stock returns. We first decompose (***OIB***) into strategy-specific components by regressing ***OIB*** on the contemporaneous sentiment scores of TA, FA, OS, and NS messages. The strategy-specific components, OIB^{TA}, OIB^{FA}, OIB^{OS} and OIB^{NS}, are the corresponding fitted values, and OIB^{Resid} is the residual. We then estimate the following return-predictability regression:

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

The control variables ($\boldsymbol{X}_{i,t}$) are the same 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	Return$_{i,t+1}$ (%)	
	(1)	(2)
OIB$_{i,t}^{BJZZ,TA}$	-3.557***	
	[-3.17]	
OIB$_{i,t}^{BJZZ,FA}$	5.555***	
	[2.98]	
OIB$_{i,t}^{BJZZ,OS}$	-6.884***	
	[-4.58]	
OIB$_{i,t}^{BJZZ,NS}$	-0.042	
	[-0.02]	
OIB$_{i,t}^{BJZZ,Resid}$	0.173***	
	[13.62]	
OIB$_{i,t}^{BHJOS,TA}$		-2.524***
		[-3.15]
OIB$_{i,t}^{BHJOS,FA}$		4.255***
		[2.99]
OIB$_{i,t}^{BHJOS,OS}$		-5.428***
		[-4.57]
OIB$_{i,t}^{BHJOS,NS}$		-0.038
		[-0.03]
OIB$_{i,t}^{BHJOS,Resid}$		0.243***
		[18.25]
Attention$_{i,t}$	-0.056***	-0.056***
	[-5.50]	[-5.55]
Stock Characteristics	Yes	Yes
Lagged Returns	Yes	Yes
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 regressions examining the contemporaneous relationship between retail investor herding on Robinhood and retail sentiment revealed by StockTwits messages, categorized by strategy type (TA, FA, OS, and NS). The regression is estimated at the stock-day level, with the following specification:

$$\textit{RH\_Herd}_{i,t} = \beta_1 \textit{Sentiment}_{i,t}^{type} + \beta_2 \textit{Attention}_{i,t} + \gamma \boldsymbol{X}_{i,t} + \boldsymbol{\delta}_t + \epsilon_{i,t}, \quad type \in \{TA, FA, OS, NS\}.$$

$\textit{RH\_Herd}_{i,t}$ is defined as an indicator equal to one if stock i is among the top ten stocks ranked by the daily percentage increase in Robinhood users holding the stock, provided that at least 100 users held stock i at the end of day $t-1$. Robinhood user account data, sourced from RobinTrack, span May 2018 through August 2020. The control variables are defined as in Table 4. All regressions include trading day fixed effects. Standard errors are clustered by stock and trading day, with corresponding t-statistics 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 (TA-BERT) for our classification task. We compare this method against two alternative approaches: (i) fine-tuning a BERT model on users' self-declared investment approaches and (ii) 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 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 from this model. The results show that this approach fails to reliably classify strategies at the message level. For example, the top-left panel shows the distribution of Technical Analysis (TA) probability for messages from self-declared technical investors. The predicted probability of TA 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 TA-BERT model against the self-declared model using a manually labeled dataset. Figure A.2 displays the TA probability for 500 human-classified messages. The top panels show that TA-BERT provides a clear classification pattern: most human-labeled technical messages receive a high TA probability, 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 differentiate between the two message types. While the average TA probability 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 zero. Likewise, messages from most fundamental investors receive a BoW fundamental score near zero.

Taken together, our analyses highlight the effectiveness of large language models in generating high-quality training data and demonstrate the ability of smaller, fine-tuned models to learn efficiently from such data for specific classification tasks.

A.2. Additional Tables and Figures

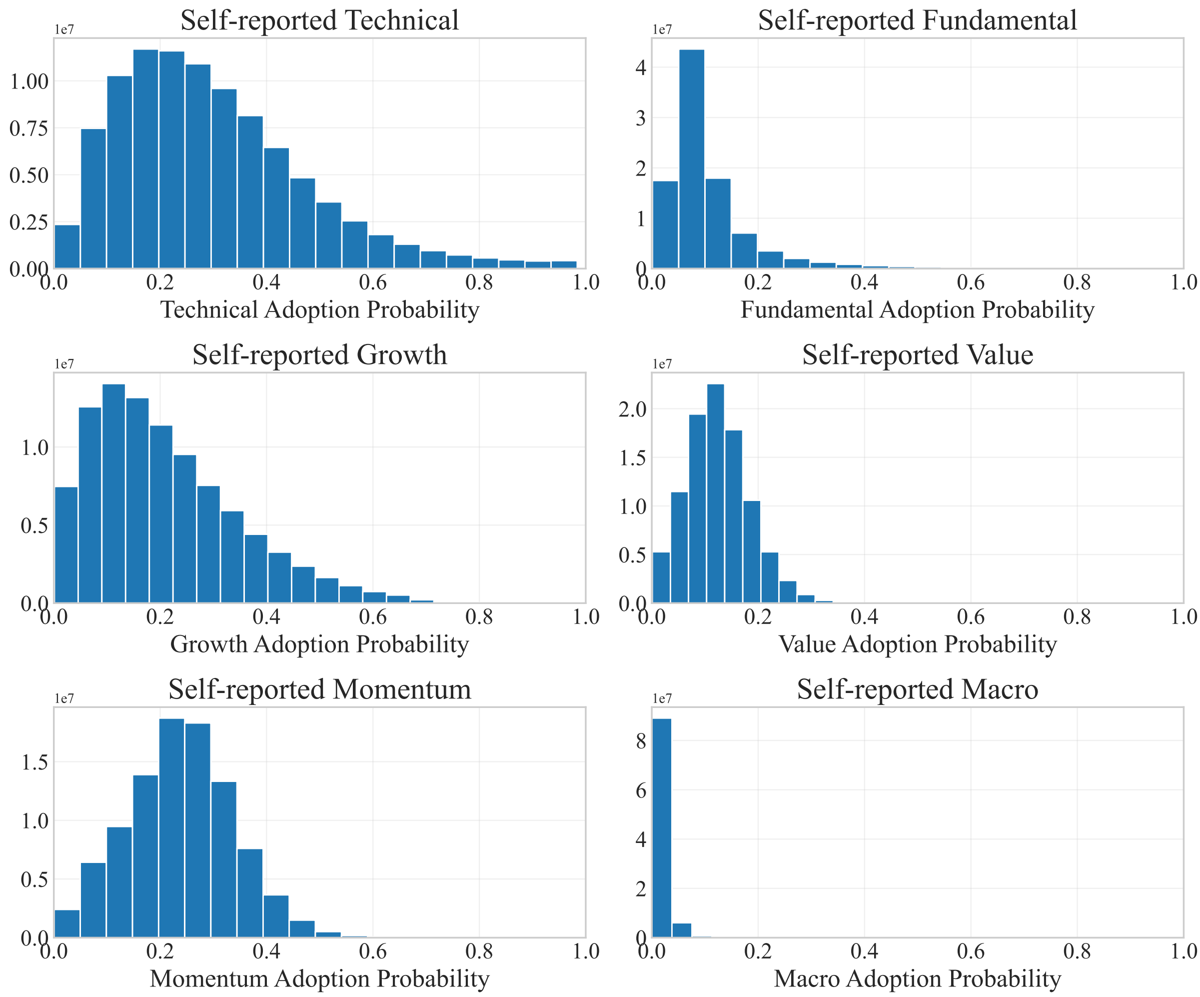

Fig. A.1. **Strategy Classification Using Self-Declared Approach and BERT**

This figure presents the distribution of investment approach classification probabilities from a BERT model fine-tuned on a random sample of messages combined with users' self-declared investment approaches. The model assigns each message a probability of belonging to one of six distinct investment approaches. For each self-declared approach, the histogram reports the frequency distribution of messages across levels of BERT-classified probabilities that a message is identified as belonging to the same approach.

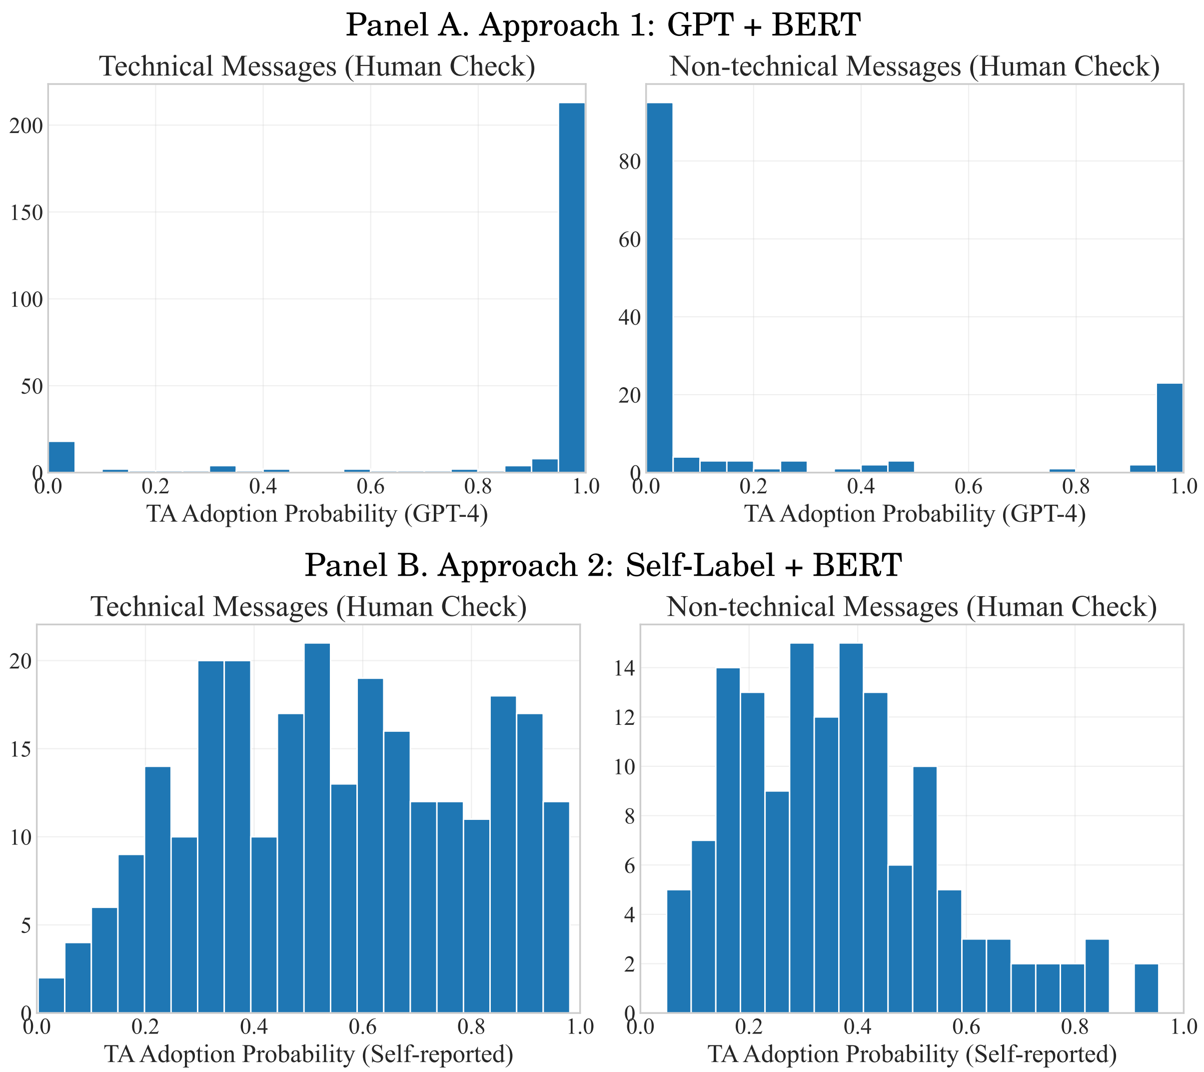

Fig. A.2. **Evaluating LLM-Based Strategy Classification Methods Against a Human-Labeled Benchmark**

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 with users' self-declared investment approaches. 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 probabilities that each 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.

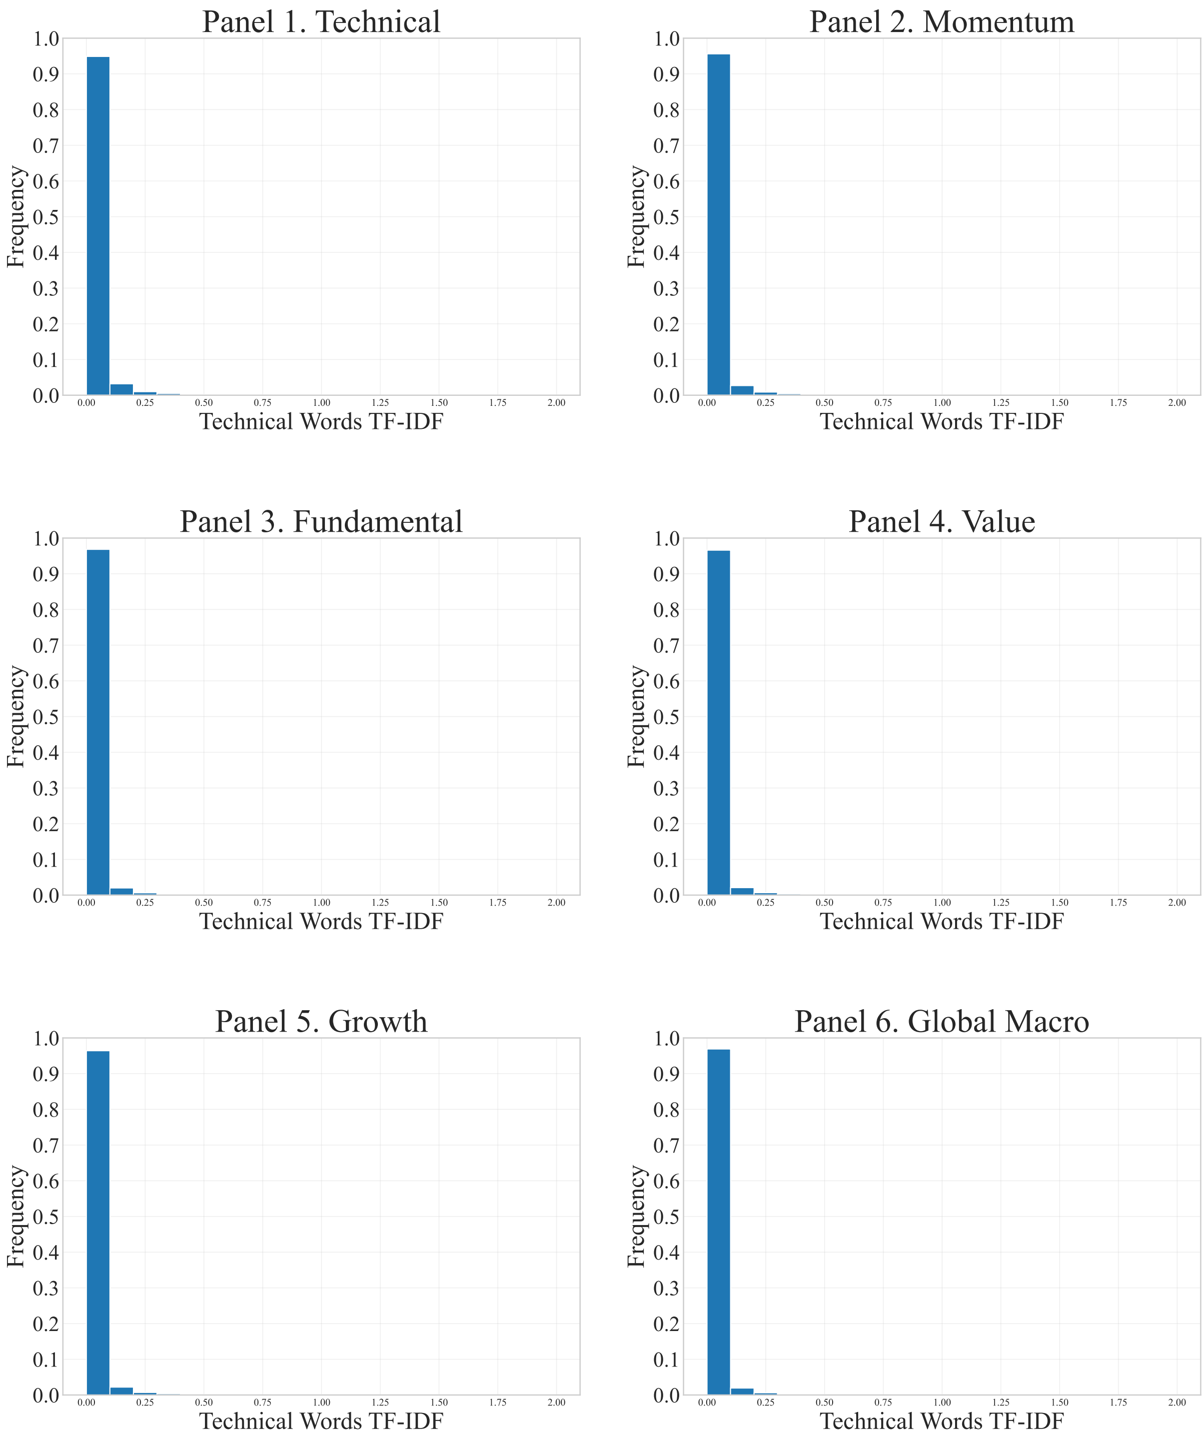

Fig. A.3. **TF-IDF Scores of Technical Analysis (TA) Words Across Self-Declared Investment Approaches**

This figure shows the distribution of TF-IDF scores for Technical Analysis (TA) words at the message level. The TA-word dictionary is obtained from Cookson and Niessner (2020). Each panel presents the frequency distribution of messages across levels of TA-word TF-IDF scores within each self-declared investment approach.

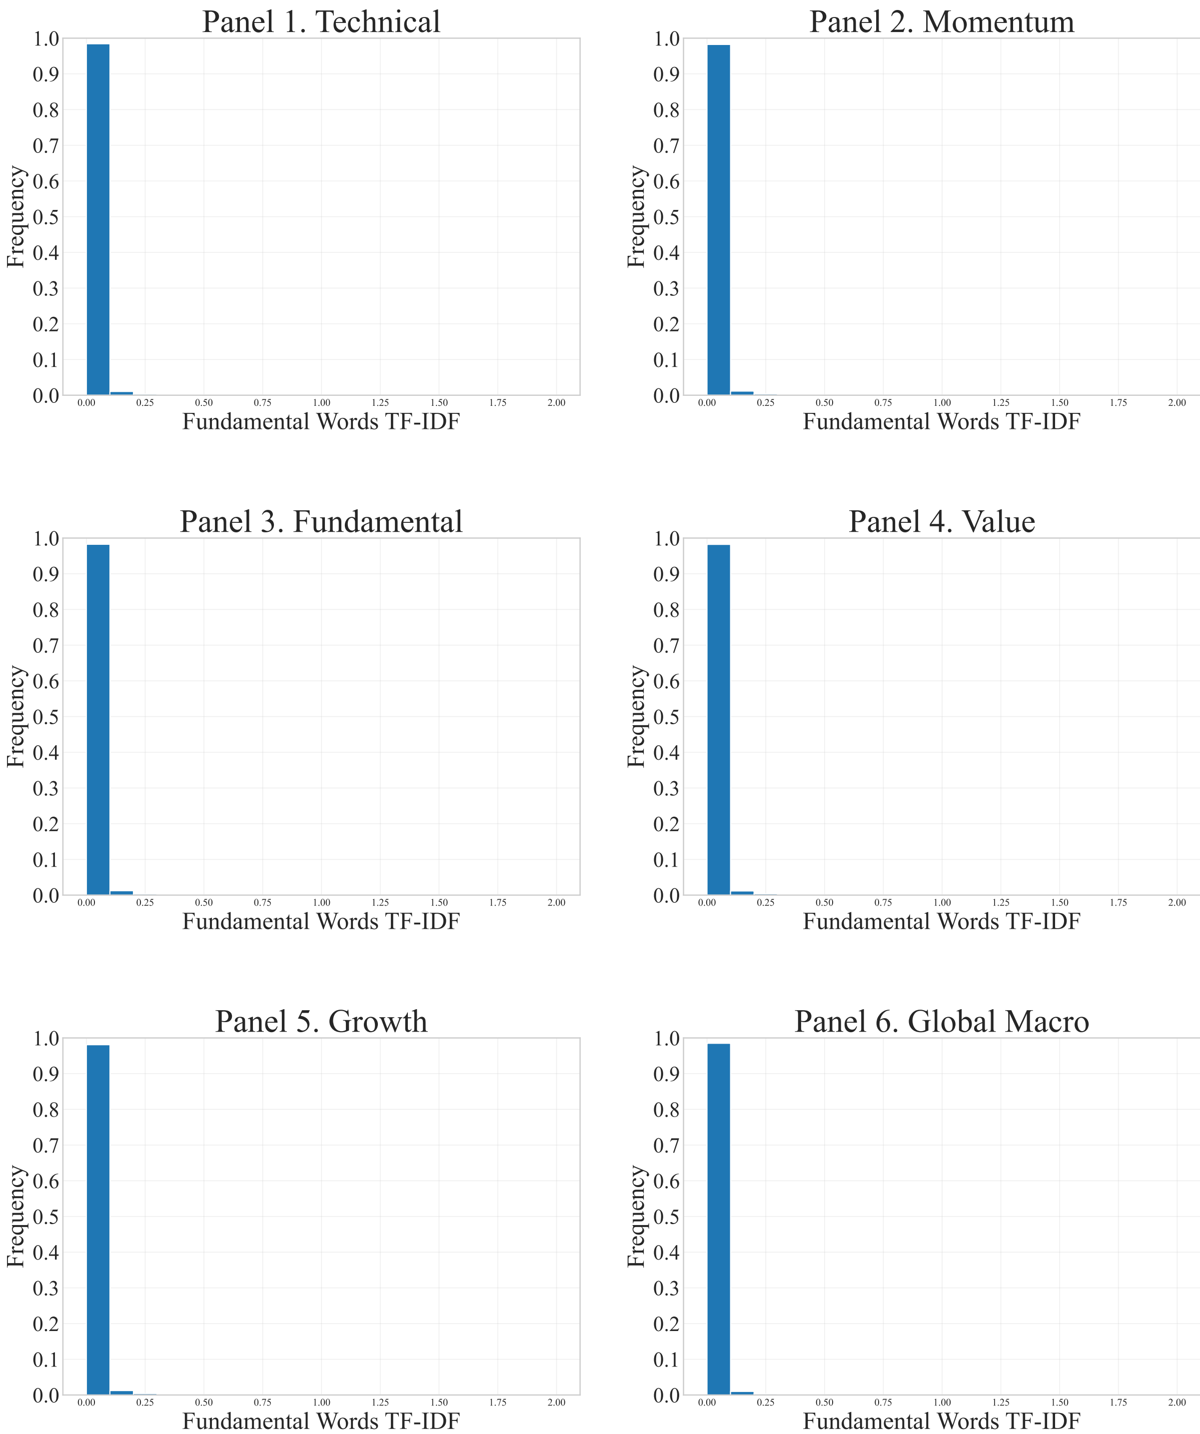

Fig. A.4. **TF-IDF Scores of Fundamental Analysis (FA) Words Across Self-Declared Investment Approaches**

This figure shows the distribution of TF-IDF scores for Fundamental Analysis (FA) words at the message level. The FA-word dictionary is obtained from Cookson and Niessner (2020). Each panel presents the frequency distribution of messages across levels of FA-word TF-IDF scores within each 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, option 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, option strategy
3	Watching $AAPL expiring 620 puts, tide could turn fast : $1.25 x 1.39:	AAPL	1	event-driven strategy, option 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, option 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	$DDD as a long term options play on the Jan 2015 Contract, Sell $40 Put, Buy $50 Call, Sell $55 Call, net credit apx $3. $$	DDD	1	option strategy
10	$AEO Rolled to June 17 15P/14C inverted strangle for $0.08DB ($1.23CR total). More time, more extrinsic value, reduce delta risk.	AEO	1	Options Trading Strategy
11	"@tunwang: $META huge earnings on mobile, way to go. will get back to above $50.?" Where it should belong higher. Bullish.	META	0	
12	$META if you dumped below $50.... good. burn with the rest pussbags	META	0	
13	$TTWO HUGE block trades: 131355 shares traded - $17.91 @ 3pm yesterday & 45900 shares traded -$17.91 @ 07:50:08 today.	TTWO	0	
14	$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	
15	$HK WTI Crude down to $94.98 & Brent down to $106.69. Could be impacking Halcon Resouces Corp.	HK	0	
16	$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**
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 by LLMs into strategy type Technical Analysis (TA), Fundamental Analysis (FA), or Other Strategy (OS).
Technical Investor$_j$	A dummy variable equal to one if investor j's self-reported investment approach is "Technical" or "Momentum".
Swing or Day Trader$_j$	A dummy variable equal to one if investor j's self-reported investment horizon is "Swing Trader" or "Day Trader".
Long-Term Investor$_j$	A dummy variable equal to one if investor j's self-reported investment horizon is "Long-Term Investor".
Professional$_j$	A dummy variable equal to one if investor j's self-reported investment experience is "Professional".
Technical$^{TF-IDF}_{i,j,t,n}$	The TF-IDF score for Technical Analysis words in a given message.
Fundamental$^{TF-IDF}_{i,j,t,n}$	The TF-IDF score for Fundamental Analysis keywords in a given message.
I(News Coverage)$_{i,t}$	Firm-level daily indicators for public news releases such as earnings news, analyst-related news (e.g., recommendations, price targets), and other business news (e.g., credit ratings, labor issues).
High (Low) Performance$^{type}_{j,t-1}$	Indicators capturing the performance of a hypothetical trading strategy following investor j's strategy-specific sentiments, which is calculated as the average subsequent 5-day stock returns across messages posted by investor j in the prior three months. High (Low) Performance is defined as an indicator variable for investors whose strategy-specific performance falls in the top (bottom) quartile (i.e., above the 75th or below the 25th percentile) in the cross-sectional distribution.
High (Low) Likes$^{type}_{j,t-1}$	Indicators capturing social feedback to investor j's adoption on each strategy type, which is measured as the total number of likes on strategy-specific messages posted by investor j in the prior three months. High (Low) Likes is defined as an indicator variable for investors whose strategy-specific like count falls in the top (bottom) quartile in the cross-sectional distribution.
Number of Words$_{i,j,t,n}$	Total number of words in a given message.
Sentiment$_{i,t}$	The difference between bullish and bearish message counts across all investors for stock i on day t, normalized by their sum (see, Cookson and Niessner, 2020).
Sentiment$^{TA}_{i,t}$	Sentiment calculated using messages related to Technical Analysis (TA).
Sentiment$^{FA}_{i,t}$	Sentiment calculated using messages related to Fundamental Analysis (FA).
Sentiment$^{OS}_{i,t}$	Sentiment calculated using messages related to Other Strategies (OS).
Sentiment$^{NS}_{i,t}$	Sentiment calculated using messages that are not related to TA, FA, or OS.
Attention$_{i,t}$	StockTwits users' 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 (see, Cookson et al., 2024a).
OIB$_{i,t}$	Retail marketable order imbalance for stock i on day t, constructed from the TAQ data following methodologies proposed in Boehmer et al. (2021) (BJZZ) or Barber et al. (2023a) (BHJOS).
RH_Herd$_{i,t}$	An indicator variable equal to one if stock i is among the top ten stocks on day t ranked by the daily percentage increase in Robinhood users, provided at least 100 users holding stock i at the end of day $t-1$ (see, Barber et al., 2022).
MAX$_{i,t}$	The maximum one-day return in the prior month.
Abnormal Turnover$_{i,t}$	A measure of abnormal trading volume, $\log(Turnover_{i,t}) - \log(\frac{1}{4}\sum_{h=1}^{4} Turnover_{i,t-h})$.
Abnormal News Volume$_{i,t}$	A measure of abnormal volume of news articles, sourced from RavenPack, $\log(1 + \#News_{i,t}) - \log(1 + \frac{1}{4}\sum_{h=1}^{4} \#News_{i,t-h})$.
Market Capitalization$_{i,t}$	Market capitalization, calculated as price times shares outstanding.
Book-to-Market$_{i,t}$	Ratio of book value to market value.
Asset Growth$_{i,t}$	Growth rate of total assets.
Gross Profit-to-Asset$_{i,t}$	Ratio of gross profit to total assets.
Analyst Coverage$_{i,t}$	Number of IBES equity analysts covering stock i.
Institutional Ownership$_{i,t}$	Fraction of shares outstanding held by 13F institutional investors.

Table A.3: **Strategy Sentiment and Future Stock Returns: Fama-MacBeth Regressions**

This table reports estimation results from Fama-MacBeth regressions of future stock returns with strategy-specific sentiments. The control variables are the same as in Table 4. 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)
$\text{Sentiment}^{TA}_{i,t}$	-0.047***	-0.055***	-0.016
	[-3.51]	[-4.35]	[-1.31]
$\text{Sentiment}^{FA}_{i,t}$	0.042***	-0.012	-0.011
	[3.36]	[-1.00]	[-0.93]
$\text{Sentiment}^{OS}_{i,t}$	-0.064***	-0.032**	-0.034***
	[-4.41]	[-2.27]	[-2.61]
$\text{Sentiment}^{NS}_{i,t}$	-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: **Strategy Sentiment and Future Stock Returns (DGTW-adjusted)**

This table reports estimation results from predictive regressions of future DGTW-adjusted returns (Daniel et al., 1997) with strategy-specific sentiments. The control variables are the same as in Table 4. All specifications 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.

	$\text{DGTW}_{i,t+1\to t+5}$ (%)	$\text{DGTW}_{i,t+6\to t+10}$ (%)	$\text{DGTW}_{i,t+11\to t+15}$ (%)
	(1)	(2)	(3)
$\text{Sentiment}^{TA}_{i,t}$	-0.059***	-0.059***	-0.024*
	[-4.17]	[-4.37]	[-1.81]
$\text{Sentiment}^{FA}_{i,t}$	0.048***	0.000	-0.002
	[3.97]	[0.04]	[-0.18]
$\text{Sentiment}^{OS}_{i,t}$	-0.063***	-0.031**	-0.038***
	[-4.10]	[-2.16]	[-2.70]
$\text{Sentiment}^{NS}_{i,t}$	-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: **Strategy Sentiment and Future Stock Returns: Stock-Day with at Least 10 Messages**

This table revisits the analysis of return predictability using the sample requiring stock-days to have at least 10 StockTwits messages. The control variables are the same as in Table 4. All specifications 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)
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: **Strategy Sentiment and Future Stock Returns: Alternative Sentiment Measures**

This table revisits the analysis of return predictability using an alternative sentiment measure computed based on the maximum entropy method. The control variables are the same as in Table 4. All specifications 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.

	Return$_{i,t+1}$ (%)				
	(1)	(2)	(3)	(4)	(5)
Sentiment$_{i,t}^{TA}$	-0.024***				-0.022***
	[-3.13]				[-3.18]
Sentiment$_{i,t}^{FA}$		0.010			0.016***
		[1.60]			[2.69]
Sentiment$_{i,t}^{OS}$			-0.032***		-0.031***
			[-3.91]		[-4.03]
Sentiment$_{i,t}^{NS}$				-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: **Strategy Sentiment and Future Stock Returns: Sub-period Analysis**

This table reports return predictability results for three sub-periods: 2010–2015, 2016–2019, and 2020–2023, reported in columns (1)–(3), respectively. The control variables are the same as in Table 4. All specifications 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.

	Return$_{i,t+1}$ (%)		
	(1)	(2)	(3)
Sentiment$_{i,t}^{TA}$	0.008	-0.018*	-0.027**
	[0.96]	[-1.72]	[-1.97]
Sentiment$_{i,t}^{FA}$	-0.005	0.028***	0.022**
	[-0.52]	[3.21]	[2.01]
Sentiment$_{i,t}^{OS}$	-0.005	-0.026**	-0.033***
	[-0.57]	[-2.49]	[-2.64]
Sentiment$_{i,t}^{NS}$	-0.001	-0.003	-0.002
	[-0.20]	[-0.30]	[-0.18]
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