Working Paper: March 2026

SMU
SINGAPORE MANAGEMENT UNIVERSITY

AI, Opinion Ecosystems, and Finance

David Hirshleifer
University of Southern California

Lin Peng
City University of New York, Baruch College

Qiguang Wang
Hong Kong Baptist University

Weichen Zhang
Singapore Management University

Xiaoyan Zhang
Tsinghua University



SIM KEE BOON INSTITUTE FOR FINANCIAL ECONOMICS

LEE KONG CHIAN SCHOOL OF BUSINESS

AI, Opinion Ecosystems, and Finance *

David Hirshleifer Lin Peng Qiguang Wang
Weichen Zhang Xiaoyan Zhang

First draft: July 2025. This draft: March 2026

Abstract

Generative AI use for content generation is associated with divergent outcomes on different financial social media platforms: indications of reasoning enhancement on Seeking Alpha, and of belief distortions on WallStreetBets. On Seeking Alpha, adoption is associated with information frictions. AI-assisted postings tilt toward analysis/credibility, and their sentiment positively predicts future returns. Use of AI is associated with more informative retail order flow and lower bid-ask spreads. In contrast, AI adoption on WallStreetBets follows surges in retail buying, and AI-assisted content is associated with emotionality and sentiment contagion. Such content precedes higher trading volume, greater volatility, and more lottery-like return distributions.

JEL Codes: D02, D82, D83, D84, D9, G11, G12, G14, G2, G4, G5, K22, O33
Keywords: Generative AI, Large Language Models, Social Media, Financial Discourse, Retail Investors, Information Frictions, Market Microstructure

*David Hirshleifer, Marshall School of Business, University of Southern California. Lin Peng: Zicklin School of Business, Baruch College/CUNY. Qiguang Wang: Hong Kong Baptist University. Weichen Zhang, Singapore Management University. Xiaoyan Zhang, Tsinghua University. The paper was previously circulated under the title "Social Finance in the Age of AI: A Tale of Two Platforms." We are grateful to GPTZero for providing access to their AI detection API and Alex Cui and Jonathan White for technical assistance. We appreciate helpful comments and discussions from Gioele Abbati, Jeremy Bertomeu, Tony Cookson (discussant), Xi Dong, Winston Wei Dou, Michael Farrell, Ioannis V. Floros, Rich Frankel, Hong Liu, Fangzhou Lu (discussant), Xiumin Martin, Tarun Ramadorai, Gideon Saar, Yuehua Tang, Liyan Yang, Michelle Yang (discussant), Bernard Yeung, Anastasia Zakolyukina, Yao Zeng, Dexin Zhou, Wu Zhu, seminar participants at Baruch College, the University of Wisconsin-Milwaukee, Washington University in St. Louis, Temple University, and participants at the Third Annual Conference on Capital Market Research in the Era of AI (AICM), the Wolfe Research 9th Annual Global Quantitative and Macro Investment Conference, the NBER Big Data, Artificial Intelligence, and Financial Economics Conference, and the 9th Annual Data Science in Finance Conference. Please send correspondence to Peng at lin.peng@baruch.cuny.edu.

1 Introduction

To function well, capital markets need to process information efficiently. The recent advent of powerful generative artificial intelligence (AI), particularly large language models (LLM) such as ChatGPT, has sharply reduced the costs of producing, disseminating, and processing information.

However, AI can also hinder the generation of useful knowledge. Stiglitz and Ventura-Bolet (2025) warns that AI can distort the information ecosystem by making it cheaper to produce misinformation and harder for consumers to evaluate content quality. Absent adequate safeguards—such as accountability mechanisms and content moderation—the authors argue that information quality may deteriorate (see also Acemoglu, Ozdaglar, and Siderius 2025). AI, especially when deployed by malicious actors, may also induce mistaken beliefs or impulsive behavior by shifting mood or by presenting triggers for for cognitive bias. So AI can increase the scope for fraud and disinformation in financial markets (Financial Stability Board 2024). Concerns about AI go beyond the financial context. Social media now occupies a central position in the general news diets of Americans.[1] Critics have argued that AI is degrading the functioning of social media opinion ecosystems. The phrase "AI slop" refers to low-quality, high-volume AI-generated content; in recognition of this phenomenon, Merriam-Webster dubbed "slop" its 2025 Word of the Year (Merriam-Webster 2025).

Synthesizing these perspectives and drawing upon both behavioral ideas and rational models, we hypothesize that the use of AI tools on digital platforms reduce the cost of producing both useful and misleading content. The net effect on the opinion ecosystem of a social media platform is therefore ambiguous. On the one hand, this technological shift can enhance the efficiency of information processing and transmission, which we call *reason enhancement*; on the other hand, this shift can degrade the ecosystem by increasing the prevalence of misinformation and emotional triggers, which we call *belief distortion*. We hypothesize that the balance between these effects depends critically upon a platform's institu-

[1] In 2025, 38% of U.S. adults say they regularly get news on Facebook and 35% on YouTube; the corresponding shares are 20% for Instagram, 20% for TikTok, 12% for X (formerly Twitter), and 9% for Reddit (Pew Research Center 2025).

tional and behavioral environment, including its governance, content moderation policies, accountability mechanisms, and the sophistication of its user base.

To test this hypothesis, we examine AI adoption in financial discourse on social-media platforms. We consider three types of questions. The first concerns selection into AI use. When do content creators choose to use AI? Is the pattern of adoption consistent with authors responding to costs of accessing and processing useful information, or with crafting promotional narratives during speculative episodes? The second concerns rhetorical format. How does AI shape the rhetorical form of financial communication, such as the use of logical argumentation, use of credible data and sources, and emotional persuasion? Third concerns downstream outcomes. How do these differences translate into user reactions and asset pricing patterns?

To measure the likelihood that a textual item is AI-generated, we use GPTZero, a leading AI content detector. GPTZero has been validated in several studies in financial and other contexts,[2] and we perform our own further validations. The reason that detection tools such as GPTZero can be effective is that AI-generated text to date tends to follow regular and predictable linguistic patterns that differ from those of human writing and its idiosyncrasies. To evaluate the effectiveness of GPTZero in detecting AI-generated discourse specifically in the context of financial social media, we construct a custom ground truth test sample. This sample consists of (1) human-written articles from SA and WSB randomly drawn from the pre-ChatGPT period and (2) LLM-generated social media content prompted with pre-ChatGPT *Wall Street Journal* news to mimic SA and WSB styles. On this dataset, GPTZero achieved an accuracy of 98% and a 0.99 F1 score, outperforming the other open-source detectors that we evaluated.[3] Our sample covers the first two years after the launch of ChatGPT, when AI tools

[2] Blankespoor, deHaan, and Li (2025) confirm the ability of GPTZero to detect AI-generated content in financial reports. Jabarian and Imas (2025) evaluate leading commercial detectors and report a false positive rate for GPTZero of 0.007 and a false negative rate between 0.002 and 0.003 on general texts across models and genres. Additional evidence on the performance of GPTZero in contexts such as news and social media is provided by Dugan et al. (2024) and Latona et al. (2024).

[3] The F1 score is the harmonic mean of precision and recall, where precision is the fraction of items flagged as AI-generated that are truly AI-generated, and recall (also called sensitivity or the true-positive rate) is the fraction of truly AI-generated items that the detector correctly flags (Christen, Hand, and Kirielle, 2023). The F1 score ranges from 0 to 1, where 1 indicates perfect performance. AI detector performance may shift as language

functioned predominantly as writing aids rather than as instruments of high-level reasoning and independent analysis. Our measure, which applies GPTZero to the final versions of written content, is therefore a reasonable proxy for AI involvement during this period.

We collect articles and messages on U.S. publicly listed firms from June 2022 to December 2024, spanning the November 30, 2022 introduction of ChatGPT, across two influential social media platforms: SeekingAlpha (SA) and Reddit's r/WallStreetBets (WSB). Seeking Alpha is by subscription only and operates under strong editorial oversight: articles are reviewed prior to publication, and authors are financially compensated for producing original investment analysis. Its user base consists largely of semi-professionals with investment expertise (see, for example, Chen et al. (2014), Farrell et al. (2022), Dim (2025)). In contrast, WallStreetBets is a decentralized Reddit community with no pre-publication review and far weaker oversight. Messages on WallStreetBets sometimes highlight screenshots of high-risk trades as well as memes and GIFs. Its users are generally less sophisticated and more prone to behavioral biases and peer effects. Together, these two platforms provide a natural laboratory for assessing how differences in platform governance and user sophistication shape both the use of AI and its economic consequences.

Applying GPTZero, we document rising AI adoption on both platforms (Figure 1). We then examine how, when, and with what consequences market participants adopt AI across these platforms. There are four key findings.

First, using Seeking Alpha author biographies, we find that AI adoption in 2023–2024 is more common among authors inferred to be non-native English speakers, research-oriented contributors, and MBA-trained generalists. This is consistent with early adopters using AI to ease language barriers and to complement analytical thinking. Authors are also more likely to adopt AI when they are less familiar with the firm and when information is sparse, as proxied by low analyst coverage or limited news availability. This evidence is consistent with AI adoption on Seeking Alpha reducing research frictions.

models and text-modification tools, such as "humanizers" designed to make AI-generated content appear more human, advance. Nonetheless, our validation results indicate that GPTZero provides a reliable and informative measure of AI usage in financial social media content during our sample period.

In contrast, on WallStreetBets, adoption is unrelated to these information-cost proxies. Instead, adoption is 44–55% higher following periods of heavy retail buying. This suggests that AI use on WSB is associated with speculative dynamics rather than to rational information processing.

Second, using OpenAI's GPT-5.1, we decompose each article's rhetoric into shares of *logos* (logic and evidence), *ethos* (credibility and authority), and *pathos* (emotional appeals and narrative framing), based on Aristotle's three modes of persuasion. AI-generated articles on Seeking Alpha are dominated by *logos* and *ethos*, consistent with platform norms that emphasize analytical rigor and the credibility of sources. On WallStreetBets, AI-generated messages place substantially greater weight on *pathos*, featuring more emotionally charged language and narrative framing.[4]

Third, on WSB, the valence of AI messages more strongly predicts the valence of subsequent comments than does the valence of human-authored messages, suggesting that AI content amplifies sentiment contagion. A possible interpretation of this finding is that AI is used to refine emotional or simplistic arguments that are contagious. The effect is absent on SA. This contrast is consistent with general evidence in persuasion research that affective appeals tend to generate more *message-congruent outcomes*, that is, responses that move in the direction implied by the message, than do cognitive appeals (Ng, See, and Cheung 2025).

Fourth, the market outcomes associated with the presence of AI messages differ sharply across platforms. On SA, the sentiment of AI-generated articles positively predicts future stock returns, with a highly bullish sentiment corresponding to a positive one-week-ahead return of 35 basis points and no evidence of subsequent reversal. In contrast, sentiment of human-generated articles on SA does not exhibit such predictive power. Moreover, SA-

[4]Aristotle's rhetorical triad is consistent with modern theories of persuasion and cognition. For example, Petty and Cacioppo (1986) suggests that persuasion can occur by means of either a "central" route (logical argumentation) or a "peripheral" route (source credibility and emotional cues). More broadly, a recent meta-analysis of 133 samples across 22 countries finds that the relative effectiveness of affective versus cognitive appeals is moderated by cultural context: affective appeals are more persuasive in collectivistic settings, whereas the two are equally effective in individualistic ones (Ng, See, and Cheung 2025). This suggests that the environment in which a message is received shapes which rhetorical mode is most influential—a pattern consistent with our finding that platform characteristics moderate the rhetorical style AI amplifies.

based AI content is associated with more informative retail order flows, as measured by the predictive power of retail net buys for future returns.[5] Additionally, stocks with AI-generated content have lower bid-ask spreads over the five days following publication. This evidence is consistent with fundamental relevant content helping resolve uncertainty and reducing the adverse selection cost faced by liquidity providers (see, e.g., Glosten and Milgrom 1985; Glosten 1989; Glosten and Harris 1988; Huang and Stoll 1997; Stoll 1989).

In contrast, AI content on WSB is associated with speculative trading and tends to be followed by extreme price run-ups. AI messages are on average followed by surges in abnormal trading volume and significantly higher volatility over the next five trading days. Moreover, this content strongly predicts lottery-like outcomes (defined as rare, extreme one-day price run-ups). Furthermore, AI content on WSB does not enhance the informativeness of retail order flow.

Exploiting Seeking Alpha's imposition of a no-AI policy in late 2023 as a quasi-natural experiment, we find that the policy altered how AI is used rather than eliminating AI use altogether. Following enforcement, wholesale AI use declined sharply, while more selective, complementary forms of AI assistance persisted. There is also evidence that AI-assisted content that remained after the ban became more informative, indicated by AI article sentiment more strongly aligned with subsequent returns. This suggests that discouraging low-effort use of AI (for example, generating a post using AI without critical oversight) can improve the average informational quality of AI-assisted analysis.

To ensure the robustness of our tests, we include stock and day fixed effects in all key specifications. For article-level tests, we further control for author characteristics or include author fixed effects. Our results are also robust to excluding days around earnings announcements and to controlling for the anticipation of near-term earnings releases. These controls address some possible concerns about unobserved confounders. Also, the results remain

[5]The evidence is consistent with AI content on SA conveying value-relevant information that is not yet fully incorporated into prices at the publication date. As investors subsequently adopt this information, either by paying attention to the AI article itself or via related post-publication news, prices on average gradually adjust. Consequently, SA AI views predict future stock performance and coincides with more informative retail order flows.

stable across alternative WallStreetBets samples with different data filters. Nevertheless, our tests are correlational rather than causal; we do not fully rule out the possibility that AI is associated with reason enhancement and belief distortion activities and outcomes without actually causing those outcomes.

However, the key alternative is that under some conditions users apply AI to generate reason enhancing content even though AI does not actually complement the generation of such content, and under other conditions apply AI to to generate belief distorting content even though AI does not actually complement the generation of such content. This is possible but not very plausible. Although not proven, the causal interpretation seems much more reasonable.

Specifically, the evidence is consistent with the fact that AI is a general-purpose tool that can be adapted to the objectives of its users and the norms of the environment in which it is deployed. AI can be deployed to sharpen analytical arguments or to amplifying emotional narratives. The consistent contrasts across the two platforms in adoption triggers, rhetoric, and market correlations suggest either that users are trying but failing to use AI for purposes relating to reason enhancement and belief distortion, or, more plausibly, that they are trying and to some extent succeeding in their endeavors.

This paper makes several contributions to the recent literature on AI use in financial markets. First, we are among the first to empirically quantify the moderating role of platform governance and user characteristics in shaping the relationship between AI deployment and market outcomes. Theory predicts that lower information costs will enhance the informational efficiency of markets (e.g., Grossman and Stiglitz 1980; Diamond and Verrecchia 1981); there is evidence consistent with this prediction (Bertomeu et al. 2025). However, concerns have been raised that AI induces misperceptions (Stiglitz and Ventura-Bolet, 2025; Acemoglu, Ozdaglar, and Siderius, 2025). Moreover, the persuasive capabilities of AI can amplify behavioral biases such as attention-induced trading (Barber et al. 2022) and social transmission bias (Hirshleifer 2020; Han, Hirshleifer, and Walden 2022) and echo chambers (Cookson, Engelberg, and Mullins 2023). Our cross-platform analysis addresses this tension by examining

how the same technology is associated with divergent outcomes depending on the institutional setting.

The finding that socially transmitted ideas are linked to retail financial behavior also fits the dual-use view of technology: tools that lower access costs can improve market functioning in high-accountability settings, but can also be used in socially dysfunctional ways when deployed without effective safeguards.[6]

Second, where recent papers have primarily focused on the consequences of AI for productivity (e.g., Noy and Zhang 2023; Chang et al. 2023; Cao et al. 2024; Bertomeu et al. 2025; Brynjolfsson, Li, and Raymond 2025), we examine the relationship of AI to financial discourse and market outcomes. Our evidence suggests that across platforms AI adoption is associated with different rhetorical strategies, user responses, and price dynamics. Taken together, these results suggest that AI may influence markets not only by lowering the cost of producing text, but by affecting belief formation and the social transmission of sentiment.

Third, this paper also broadens the scope of the literature on AI analysis of textual communication in finance by studying retail-facing online platforms rather than disclosure channels such as corporate disclosures (Blankespoor, deHaan, and Li 2025) and professional analyst reports (Bertomeu et al. 2025). Our study provides one of the first large-scale comparative analyses of the footprint of AI in the more dynamic and loosely regulated domain of public equity commentary. We document how AI reshapes the production and dissemination of arguments and narratives in settings where persuasion plays is central.

Bradshaw et al. (2025) (hereafter BMYZ) study AI deployment on Seeking Alpha and find that authors who adopt AI become more productive, as measured by publishing more articles and covering more new firms than nonadopters. A key difference between our paper and BMYZ is that we examine AI use in content generation on different platforms and find that associated outcomes diverge sharply. The evidence of BMYZ is broadly consistent with

[6]AI and related financial technologies can reduce biases in portfolio choices and improve household financial decisions (D'Acunto, Prabhala, and Rossi 2019; Rossi and Utkus 2024; Lu, Li, and Shi 2025), the advice generated by LLMs can also reflect demographic and preference biases that vary across models and prompts (Choukhmane et al. 2025; Fedyk et al. 2025). See Campbell and Ramadorai 2025 and Heimer 2026 for reviews of household financial decisions and the role of AI and related technologies.

the reason-enhancement role of AI that we find for Seeking Alpha, whereas our evidence from WallStreetBets highlights a potential dark side of AI use.

A caveat to our conclusions is that our proxy for AI usage is imperfect. It does not capture upstream activities such as screening stocks, conducting background analysis, or refining pre-existing arguments. Furthermore, heavily rewritten AI drafts are unlikely to be detected. In later samples that might be employed by future studies, these challenges are likely to intensify as AI capabilities advance and AI-assisted content becomes harder to identify (though detectors are likely to improve as well). Our results should accordingly be interpreted as characterizing detectable AI use during this early adoption period and the corresponding market dynamics. Even so, the core insight of our study—that the institutional environment shapes whether AI improves or degrades the quality of financial discourse—likely extends beyond the specific detection technology or sample period. As AI tools grow more powerful, both reason enhancement and belief distortion are likely to become more pronounced, and the societal challenges they pose will only deepen as AI use becomes more sophisticated and harder to observe.

A further caveat to our conclusions is that our cross-platform comparison capture multiple institutional differences—governance, content moderation, user sophistication, and compensation structure. It is not possible to fully disentangle the effects of these differences, though we have argued that these differences offer predictions in the same direction across platforms. Furthermore, the imposition by Seeking Alpha of a no-AI policy in late 2023 provides a quasi-natural experiment that speaks directly to the role of platform governance. Following this policy change, detectable AI use declined, usage shifted toward more selective and likely higher-level forms of AI assistance, and remaining AI content became more informative. This finding indicates that platform regulation can meaningfully shape how AI is deployed within an opinion ecosystem.

2 Hypotheses

In markets with information frictions, tools such as AI that lower production costs or alter how readers process content can affect discussion patterns and market outcomes. Motivated in part by the model of Stiglitz and Ventura-Bolet (2025), we argue that the magnitude and direction of these effects are critically shaped by the institutional and behavioral environment of a platform. The sharp contrast between Seeking Alpha and WallStreetBets in their governance and user base provides a natural setting to test this idea. We first outline these institutional differences and then develop hypotheses regarding AI adoption motives, user responses, and associated capital market outcomes.

2.1 Institutional Background

Seeking Alpha and WallStreetBets differ fundamentally in governance, user sophistication, and content norms. SA fosters reputation-building through in-depth analysis, whereas WSB rewards engagement through memes and high-risk/high-gain narratives.

Seeking Alpha, founded in 2004, is a platform for its crowd-sourced and curated equity research.[7] Its contributors include a network of stock analysts, traders, economists, academics, financial advisors and industry experts. As of 2021, the platform had more than 17,000 individual contributors and over 350 contributing firms. An in-house team of over 60 editors reviews every article submission, providing iterative feedback to authors until the content meets the platform's standards—a process that typically takes 10-24 hours for weekday submissions. This moderation extends to the comments section, which is actively monitored to maintain substantive discussion. Seeking Alpha has more than 10 million registered users and over 17 million monthly visitors, with an average visit duration double that of Morningstar and Yahoo Finance, and four times that of The Economist, Barron's, and the Wall Street Journal.

This structured governance of Seeking Alpha is reinforced by content guidelines for con-

[7] The information presented in this subsection for Seeking Alpha is based on content from its official website, for example, https://static.seekingalpha.com/uploads/pdf_income/sa_media_kit_2021_final.pdf.

tributors. The platform encourages “investment-oriented”, “original”, and “fundamental analysis”, and explicitly prioritizes articles that help an investor in their decision making. Furthermore, authors must disclose any business relationships with or personal holdings in the companies they cover. Seeking Alpha writers are paid a fixed amount for each article and a variable amount based on views by paid subscribers. Despite allowing pseudonyms, the platform requires private identity disclosure and prohibits users from posting under multiple accounts.[8]

Seeking Alpha’s user base also reflects this professional focus. The average retail user is a 46-year-old male (80% of the user base) with a reported household income of $321,000 and $1.5 million in investable assets. Behaviorally, they are active investors, conducting an average of 11.6 transactions annually. Moreover, a large majority are considered "affluencers" (80%) and act as the lead financial decision-maker in their household (76%). The platform’s average professional client manages over $165 million in assets (AUM) and possesses deep industry experience, with 93% having worked in finance for at least a decade. Their practice is typically focused on an elite clientele, as 80% of these professionals managing assets for high-net-worth individuals.

In contrast, WallStreetBets (WSB) is a decentralized Reddit community, created on January 2012, known for its meme-centric discussions with a particular focus on highly speculative trading strategies. As the largest finance-related subreddit with over 13 million subscribers, its substantial growth is largely attributable to the GameStop (GME) short squeeze event, during which the forum expanded from 500,000 users in July 2018 to 10.7 million by June 2021 (Bradley et al. 2023).[9] The conventional view is that WallStreetBets’ user base consists of a less affluent and less sophisticated retail investor demographic.

WSB messages appear instantly without any pre-publication review, and there is only limited post-publication moderation for spam and abuse. Messages often highlight high-

[8] See https://about.seekingalpha.com/policy-on-pseudonymous-analysts.

[9] New members (those who joined WSB in early 2021) tended to emphasize coordinated buy-and-hold strategies for a handful of meme stocks, with little regard to the company’s fundamentals. Indeed, WSB users have been characterized as a “mob of uninformed, unsophisticated retail traders” (John Coffee as quoted in Detrixhe 2021). Also see Asarch (2021) for the anecdotal evidence on new WSB members.

risk, speculative trades. In line with its speculative culture, WallStreetBets encourages users to share screenshots of high-risk holdings and trades through designated "YOLO" and "Gains/Losses" tags (known as 'flairs' on Reddit).[10] The highlighting of trading gains versus losses is likely to trigger reference-dependent thinking (Shefrin and Statman (1985)) , a form of imperfectly rational cognition that has been heavily documented in behavioral economics and finance. There is no direct monetary compensation for posting on WSB. Instead, the reward is social capital, which includes karma points (a quantitative score reflecting a user's reputation and contributions, earned through upvotes from other members), and collective recognition and sense of identity as a community member. WSB permits anonymity and there are no restrictions on using multiple accounts.

Taken together, these structural differences in governance, content guidelines, and user sophistication and incentives create distinct environments for AI deployment and its relationship with markets.

2.2 Adoption of AI in Content Creation

AI can be used to produce both meaningful and misleading content. Under what we call the *Reason Enhancement Hypothesis*, AI strengthens a platform's role in knowledge discovery by helping creators extract, synthesize, and communicate value-relevant signals. Under the *Belief Distortion Hypothesis*, AI contributes instead to the generation of vivid, affectively charged, or persuasive narratives that diverge from fundamental information, thereby degrading the informational environment.

Based on this, we develop specific predictions regarding the adoption of AI in financial content creation.

Under the Reason Enhancement motive for adoption, AI complements information ac-

[10] See https://www.reddit.com/r/wallstreetbets/wiki/linkflair/. For example, users can select from a list of flairs to add to their message, where "YOLO" corresponds to trades with minimum value of risk at $10,000 in options, or $25,000 in equity, the terms "gain" and "loss" are reserved for posts that document a solid winning or losing trade, respectively, with at least $2,500 in options profit or loss or $10,000 in share profit or loss. The WSB content guidelines focus primarily on combating spam and abuse with rules such as "no crypto," "no penny stocks," "no paper trading," "no pump-and-dumps," "no market manipulation," and "no social begging." Moderators also intervene to remove messages that misuse flairs.

quisition, learning, and analytical production. Creators use AI to generate and communicate value-relevant analysis and to clarify, structure, and synthesize available information. These benefits are particularly important when authors face higher baseline costs of producing such content. Such costs are higher when content creation requires overcoming language frictions or steep learning curves, including non-native language use, limited prior familiarity with the covered firm, or facing complexity.

Importantly, the pedagogical value of AI is not limited to novice users. Research-oriented authors and generalists, such as analysts and individuals with PhDs, often need to learn about new settings, organizing dispersed signals, and translate analysis into coherent narratives. Moreover, such authors may be especially well positioned to benefit from AI assistance, as domain expertise and analytical training complement the strength of AI in rapid synthesis and structuring, allowing users to sharpen argument and to revise more efficiently.

This reason-enhancement motive is especially salient when the surrounding information environment is poor. When public information is sparse and investigation is more difficult, the ability of AI to organize, contextualize, and interpret information becomes particularly valuable. We therefore expect AI adoption to be more frequent among non-native authors, among authors engaged in analytically intensive or research-oriented tasks, and when an author covers a firm for the first time or when public news availability is low.

The second mechanism, *Belief Distortion*, arises when creators use AI to generate narratives that are systematically inaccurate, noisy or misleading. Such narratives may emerge because creators themselves hold biased or misinformed beliefs, or because creators strategically deploy inaccurate narratives to exploit viewers' behavioral biases. In either case, AI enables the production of emotionally resonant, sensational, or identity-affirming content designed to shape beliefs or mobilize attention. This can be without regard to accuracy or with deliberate intent to foster inaccuracy. Experimental evidence shows that LLMs outperform humans in persuasion tasks (Schoenegger et al. 2025) and can be turned into persuasion tools with minimal prompting (Hackenburg et al. 2025). We refer to this broad class of persuasive, non-informational uses as *persuasion engineering*, which includes both attention-

grabbing visual content (e.g., images and memes) and crafted narratives that we term *rhetorical engineering* (Matz et al. 2024; Cong et al. 2024). Such uses are especially valuable on lightly moderated platforms with less sophisticated users who are more susceptible to emotions and social influence.

We therefore hypothesize:

Hypothesis 1a *The probability of AI adoption for creating financial content is higher when authors face greater language or learning frictions, when an author covers a firm for the first time, and when the public information environment is sparse. AI adoption is also higher among research-oriented authors and analytical generalists.*

Hypothesis 1b *The probability of AI adoption for creating financial content is positively associated with recent surges in retail buying activity.*

On curated platforms such as Seeking Alpha, where editorial screening ensures baseline quality standards and users are relatively sophisticated, the incentive structure rewards well-researched fundamental analysis. This environment makes AI a natural complement to value-relevant information production, so AI adoption is more likely to reflect the reason-enhancement motive. In contrast, on lightly moderated and speculative platforms such as WallStreetBets, AI adoption is more likely to reflect information- distortion motives, arising either from creators' own biased beliefs or from attempts to shape or exploit the behavioral biases of others through emotionally charged, entertaining, or persuasive content. In this environment, AI-enabled persuasion is particularly valuable during episodes of retail-driven trading disconnected from fundamentals. We therefore expect Hypothesis 1a to be most relevant for SeekingAlpha, and Hypothesis 1b for WallStreetBets.

2.3 AI-Generated Rhetoric and User Responses

We now turn to how AI-generated content shapes user responses through rhetoric, which we characterize using Aristotle's three modes of persuasion as mentioned earlier: *logos* (logic- and evidence-based argumentation), *ethos* (appeals to credibility and authority), and *pathos* (emotional appeals). On curated platforms that reward analytical rigor and the credibility of

data and sources, AI can be used to produce structured, evidence-based arguments and to strengthen credibility cues, thereby increasing *logos* and *ethos*. On lightly regulated and more speculative platforms, AI can also be used to generate vivid and emotionally resonant narratives that mobilize attention and engagement, thereby increasing *pathos*. This argument motivates our next hypothesis.

Hypothesis 2 *AI's rhetorical engineering is context dependent. On curated platforms that reward analytical rigor and credibility, AI-generated content relies more on evidence-based reasoning and on language that conveys source credibility. On lightly regulated and more speculative platforms, AI-generated content relies more on emotionally charged and attention-grabbing language.*

We next consider how users respond to AI-generated financial content. We hypothesize that if AI is used to synthesize and disseminate value-relevant information, it will support rational information processing; when it is used to craft distorted or affectively charged narratives, it will amplify behavioral biases and sentiment propagation, especially on lightly moderated platforms with less sophisticated users such as WallStreetBets.

Persuasion via the reason-enhancement mechanisms (*logos* and *ethos*) operates by organizing information and strengthening perceived source credibility, and therefore should not in general induce belief overreaction. In contrast, persuasion through *pathos* relies on emotional salience and narrative intensity, which can elicit stronger affective responses and belief overshooting. As a result, sentiment contagion is more likely to arise when AI-generated content emphasizes *pathos*, particularly on speculative platforms.

These dynamics of information degradation align with existing theories of bias in social transmission. For example, in the model of Han, Hirshleifer, and Walden (2022), the social transmission of a salient signal about a stock, such as a large positive return, can induce naïve receivers to overweight this signal in forming beliefs about future performance. Alternatively, in the model of Hirshleifer (2020), social transmission can systematically cause signals about a stock to be biased upward, which can induce overoptimism among receivers. Furthermore, there are models in which optimism versus pessimism about a stock spreads

via social network (Burnside, Eichenbaum, and Rebelo 2016, Hirshleifer 2020, Pedersen 2022). In our context, if AI enables an author to make a content more salient and persuasive, unsophisticated viewers are likely to overreact, thereby amplifying sentiment contagion and producing tighter alignment between viewer comments and the sentiment expressed in the initial message. However, we must keep in mind that contagion of sentiment (i.e., valence) could also occur by rational updating.

Based on these considerations, we propose the following.

Hypothesis 3 *AI-generated content is associated with greater sentiment contagion, where the sentiment of viewer comments more closely aligns with the sentiment of the initial content.*

To the extent that sentiment is spread by emotional or simplistic arguments, we expect Hypothesis 3 to be more applicable to WallStreetBets.

2.4 AI-Generated Content and Capital Market Outcomes

Finally, we consider how AI-generated content maps into capital market outcomes. AI deployment can enhance information efficiency by facilitating genuine information production and dissemination or can amplify speculative dynamics by reinforcing behavioral distortions.

On curated platforms such as Seeking Alpha, where editorial oversight and user sophistication channel AI toward value-relevant analysis, we expect AI deployment to help investors accurately evaluate firms. If these signals are not fully and immediately incorporated into prices, attentive investors—particularly rational retail traders facing high information costs—may earn higher abnormal returns. Consistent with evidence that aggregate retail buying predicts positive future returns (Kaniel, Saar, and Titman 2008; Barrot, Kaniel, and Sraer 2016; Boehmer et al. 2021), AI-generated content may further strengthen this informativeness by lowering information frictions for rational retail participants.

Furthermore, the reason-enhancement channel implies improvements in information discovery. By reducing uncertainty and mitigating adverse selection (see, e.g., Glosten and Milgrom 1985; Glosten 1989; Glosten and Harris 1988; Huang and Stoll 1997; Stoll 1989), AI-

generated value-relevant content should narrow bid-ask spreads. More broadly, if AI facilitates the incorporation of fundamental information into prices, it should oppose price bubbles and other deviations from intrinsic value.

In contrast, on lightly moderated and speculative platforms such as WallStreetBets, AI adoption is expected to more often derive from belief-distortion motives. This results in opposite implications for information discovery in capital markets. As discussed in Hypothesis 3, AI-generated discourse that reinforces the biases of creators or manipulates the biases of readers can heighten sentiment contagion.[11] Such content encourages noise trading (Black 1986, Barber and Odean 2000), resulting in higher speculative trading volume and greater excess volatility. Consistent with Cookson, Engelberg, and Mullins (2023), who find that echo chambers stimulate trading activity, AI deployment may intensify these dynamics by making persuasive or emotionally charged posts more salient and viral.

Distortionary AI usage can also amplify social-transmission bias (Hirshleifer 2020), in which social interactions systematically shift ideas or signals in a particular direction as they pass from person to person. The interaction of social transmission bias with short-selling constraints—common in retail-dominated markets—can increase the probability of sharp price run-ups unanchored from fundamentals (e.g., Bali et al. 2025). AI usage is then predicted to generate lottery-like return distributions characterized by occasional extreme upside movements and lower expected returns.

The effect of AI on bid-ask spreads in speculative settings is theoretically ambiguous. Noise trading may reduce spreads by reducing adverse selection (e.g., Glosten and Milgrom 1985; Kyle 1985). However, recent evidence suggests that the opposite occurs when inexperienced investors herd and create inventory risks that harm liquidity (Eaton et al., 2022). Thus, the relationship between AI content and spreads is determined by the relative strength of the adverse-selection channel versus the retail-clustering channel, making the overall outcome an empirical question.

Based on these considerations, we propose the following two hypotheses linking AI-

[11] Sentiment, measured as valence, can also be spread by rational updating. These possibilities can be distinguished by testing for price reactions and reversals.

generated content to market outcomes:

Hypothesis 4 *AI-generated content is associated with improved price discovery, as reflected in narrower bid-ask spreads, more informative retail order flows, and lower likelihood of price bubbles.*

Hypothesis 5 *AI-generated content is associated with more speculative trading, as reflected in higher subsequent abnormal trading volume and return volatility, and more positively skewed ("lottery-like") short-horizon return distributions.*

As with user responses, we expect that Hypothesis 4 will be more relevant for Seeking Alpha, where AI supports reason-enhancement motives, while Hypothesis 5 will be more relevant for WallStreetBets, where AI amplifies speculative dynamics and sentiment contagion.

3 Data

This section details the data sources and variable construction for our empirical analysis. We begin by introducing our primary data from Seeking Alpha (SA) and Reddit's WallStreetBets (WSB). We then describe our methodology for identifying AI-generated content. Finally, we define the key variables used in our tests and present summary statistics.

3.1 Social Media Financial Content

Our primary data consist of the textual content collected from two online platforms: Seeking Alpha (SA) and Reddit's r/WallStreetBets forum (WSB). The sample period spans from December 2022, immediately after the launch of ChatGPT in November 2022, to December 2024.

For the Seeking Alpha sample, we collect 83,654 articles published between December 1, 2022, and December 31, 2024, within the Analysis section, excluding earnings call transcripts and corporate presentations to focus on original analytical work. Following past literature, we further restrict the sample to include (1) single-ticker articles (Chen et al. 2014), and (2) multi-ticker articles for which the author specifies a primary ticker identifying the focal firm

(Campbell, DeAngelis, and Moon 2019). 57,581 articles meet this criterion.

For WallStreetBets, we consider all messages, including submissions (original posts) and subsequent comments (follow-up replies within a thread). The raw dataset contains 969,747 messages, with submissions accounting for 10.35% of the total. We apply two filters. First, to enable reliable detection of AI-generated content, messages must exceed 50 words, leaving 82,523 messages. Second, messages must reference a single valid stock symbol; this yields 64,848 messages.[12]

We process individual content with GPTZero Model 3.1b (Base build, January 9, 2025) with default settings. To assess whether GPTZero has substantial resolution ability for financial text, we conduct a validation study using 1,000 pre-ChatGPT human-written SA and WSB articles, along with AI-generated counterparts created from pre-ChatGPT *Wall Street Journal* news articles using multiple LLMs prompted to mimic SA and WSB (See Internet Appendix A for details). As shown in Appendix Table A2, GPTZero achieved an accuracy of 98% and an F1 score of 0.99, outperforming mainstream open-source detectors (e.g., Wu et al. 2025) in distinguishing AI from human content in this setting.[13]

Figure 1 plots the monthly fraction of AI content on each platform. On SA, adoption rises sharply after the release of ChatGPT, peaking at 13.7% in late 2023. The fraction then declines over the subsequent months, following the platform's efforts to discourage AI use, including detection using proprietary technology and the suspension or removal of authors for violations. Even so, the percentage of AI content as detected by GPTZero remained stable at approximately 3–4% through the end of 2024.[14]

On WSB, adoption was slower but steady: AI content remains under 2% until mid-2023

[12]Section 7.4 presents robustness checks that include all WSB messages, regardless of length. Due diligence posts are rare in our sample (1,617 posts, 0.17%), so we do not analyze them separately (such posts are the focus of Bradley et al. (2023)).

[13]Accuracy is the share of documents correctly classified as AI-generated or human-written. The F1 score is the harmonic mean of precision and recall, providing a single measure that balances between false positively and false negatives.

[14]This trend is consistent with the findings of Bradshaw et al. (2025). In a notable shift, Seeking Alpha introduced its own "Analyze with AI" feature around February 2025 as part of an update to its Premium subscription Advanced Charting tool. The feature uses generative AI to produce an up-to-date, easy-to-read company report. In Section 7.2, we examine how AI usage patterns of authors change following the policy shift and revisit our key results before versus after the change.

and then climbs gradually to around 7% by December 2024.[15] The slower uptake is likely associated with the short, slang-heavy, and meme-driven format of WSB, where human spontaneity is already low-cost, reducing the relative advantage of early AI tools. In comparison, Blankespoor, deHaan, and Li (2025) finds that the greatest use of AI in corporate filings is in S1 filings (for new security issuances), with rates reaching between 2.5% and 4.5% in 2024. This suggests that social media commentators adopted AI more aggressively than did corporate filers.

3.2 Variables Construction

We collect financial market data from a number of sources. Daily stock returns, trading volume, and pricing information are from the Center for Research in Security Prices (CRSP); firm-level accounting data are from Compustat; quarterly institutional ownership is from Thomson Reuters Institutional Holdings (13F) database; and analyst coverage is from Institutional Brokers' Estimate System (IBES). To construct microstructure measures, we use the Trade and Quote (TAQ) database.

We define a *piece* as an article or message. We construct variables at two levels: the individual content level and the aggregated stock-day level (summed over individual pieces for the same stock on the same day). Here, we focus on the key AI variables and defer the detailed definitions of all other commonly used variables to Appendix Table A.4.

Article-level variables GPTZero analyzes each piece and outputs the probabilities that it belongs to one of three categories: "AI-ONLY," "HUMAN-ONLY," or "MIXED" (human with AI assistance). These probabilities sum to one. For example, an article may have probabilities of 0.58, 0.02, and 0.40 for being "AI-ONLY," "MIXED," and "HUMAN-ONLY," respectively. Based on these outputs, we define a continuous AI probability score, $AIProb_{ijt}$, as the sum of the

[15] The sharp increase in the share of AI-generated content on WallStreetBets in December 2024 was largely driven by the final holiday week (December 24–31), which had lower overall messaging activity but a higher share of AI-generated content. During this week, there were a total of 540 messages, of which 75 (13.2%) are classified as AI-generated. In comparison, over the first three weeks of December, the corresponding weekly totals are 854 messages and 52.3 AI messages (6.1%).

"AI-ONLY" and "MIXED" probabilities of piece j covering stock i on day t, and use $AIProb_{ijt}$ for our piece-level analysis. In this example, the piece's *AIProb* would be 0.60.

The API also outputs a categorical label for each article, determined by the category with the highest probability. In the same example, the article would be classified as "AI-ONLY" because its probability is the highest among the three categories.[16]

Other article-level variables include length (word count); *Sentiment*, measured as the number of positive words minus the number of negative words, divided by total word count; textual complexity, measured by the proportion of complexity words (Loughran and McDonald 2024); the Fog Index, a readability measure.[17] We also include indicators for quantitative content (count of numbers) and graphical content (count of images). Viewer disagreement is measured as the standard deviation of sentiment scores across comments posted within the 10-day window beginning on the publication date, calculated using platform-specific sentiment dictionaries.

Stock-day level variables At the stock-day level, we define $AIDay_{it}^{SA} = 1$ if at least one SA article about stock i on day t is classified as "AI-ONLY" or "MIXED," and 0 otherwise. Similarly, we define $AIDay_{it}^{WSB}$ for WSB messages. These platform-specific AI indicators serve as our key measures of AI adoption in the stock-level analysis.[18]

We align publication times with market trading: articles posted before 4:00 p.m. are assigned to the same trading day, while those posted after 4:00 p.m. or on non-trading days are

[16] Appendix Table A3 presents representative examples. An AI-generated Seeking Alpha article about Nike (97.75% AI-Only probability) contains formulaic language such as "These risks, while manageable, necessitate strategic planning and operational flexibility to mitigate potential impacts on revenue and profitability." GPTZero flags mechanical transitions, task-oriented structure, and lack of creativity. In contrast, a human-written article about Microsoft (99.25% Human probability) exhibits distinctive voice: "Everyone knows ChatGPT. It is a smart coder, an intelligent writer, and an elegant Chat Bot... For the first time, search becomes a peaceful and delightful human-like dialogue." On WallStreetBets, an AI-generated message about Intel (99.99% AI-Only) displays polished structure atypical of the platform's informal register, while a human message about Adobe uses casual, first-person narration characteristic of retail investors.

[17] The Fog Index is defined as $0.4 \times \frac{\text{Words}}{\text{Sentences}} + 100 \times \frac{\text{Complex Words}}{\text{Words}}$, where complex words are have three or more syllables.

[18] Our main tests focuses on AI probability measured at the article level or at the stock–platform–day level. In additional tests (Subsection 7.1), we use sentence-level AI probability measures to characterize how AI-like content is integrated into the structure of an article.

assigned to the next trading day. Outcome variables are measured after publication, typically over $t+1$ to $t+5$: cumulative abnormal returns (CAR, cumulative buy-and-hold returns adjusted by the market return); abnormal trading volume (AVOL, mean daily log dollar volume over $t+1$ to $t+5$ minus the mean over $t-41$ to $t-11$); effective bid-ask spreads (Spread, mean daily volume-weighted effective bid-ask spread); and realized volatility (Vol, standard deviation of the daily stock returns).

3.3 Summary Statistics and Descriptive Analysis

Table 1 presents the summary statistics for our sample. Panel A reports article characteristics. SA articles are substantially longer than that of WSB (1,454 vs. 114 words), more positive in tone, more complex, and contain more quantitative elements and figures. SA articles also attract more comments (11.3), whereas WSB messages draw significantly fewer comments (3.85). These patterns align with the emphasis of SA on in-depth analysis relative to the shorter, meme-centric style of WSB. SA articles have nearly three times the AI probability (10.6% vs. 3.6%), suggesting that authors on SA are more likely to use AI as flagged by GPTZero.

Panel B summarizes author characteristics. Our sample includes 1,998 unique authors on SA and 28,563 on WSB. Of these, 457 (23%) SA authors and 19,025 (67%) WSB authors publish only once; these one-time contributors account for 0.79% of SA and 22.8% of WSB articles, respectively. On average, a SA author publishes 28.78 articles and covers 16.24 stocks across 10.24 four-digit SIC industries for our sample period of December 1, 2022 to December 31, 2024. In comparison, the average WSB author published 2.27 messages and covers 1.64 stocks across 1.60 industries. Overall, these statistics indicate that SA authors are less numerous but more prolific, covering a broader range of firms, while the vast number of WSB authors contributes less frequently on an individual basis.

Panel C reports the characteristics of the stock-day sample. SA covers more stocks than WSB (3,193 vs. 1,351), resulting a larger stock-day sample. However, for the stocks covered, WSB exhibits higher coverage intensity, generating 9.43 articles per 100 stock-days compared

to 3.59 for SA. AI days occur in 0.24% of SA stock–days and 0.18% of WSB stock–days. Firm characteristics indicate that, compared to WSB, SA coverage tilts toward smaller firms with higher book-to-market ratios, lower news coverage, and less analyst coverage. Liquidity profiles also differ: stocks discussed on SA exhibit wider spreads (0.41% vs. 0.29%). Taken together, Table 1 establishes SA and WSB as distinct ecosystems with different informational, engagement, and liquidity environments.

Appendix Table A5 describes the association between AI usage and the textual attributes of pieces. Columns 2, 4, and 6 including author fixed effects for SA and WSB pieces, respectively; columns 5 and 6 also add an indicator for initial submissions to distinguish them from replies. The result indicates that AI content is generally more positive in sentiment, higher in complexity, and harder to read (as measured by the Fog Index).[19] On SA, AI articles also contain fewer images and numbers, whereas on WSB they tend to include fewer images but more numbers—though these associations become insignificant once author fixed effects are included. Notably, WSB submissions exhibit a 6.96 percentage points (or 193% relative to the sample mean of 3.6 percentage points) higher AI probability than comments, indicating that users are more likely to rely on AI for creating original posts than for replies.

To mitigate the concern that Seeking Alpha and WallStreetBets cover different stock universes, and that cross-platform differences may therefore reflect stock selection rather than platform governance, we include stock fixed effects in all specifications to absorb time-invariant stock-level differences. We also include time fixed effects and either control for author characteristics or include author fixed effects when applicable. In addition, our market-outcome analysis uses a combined stock-day-platform panel that pools observations from both platforms, enabling comparisons within the same stock-day across platform contexts.

We next describe the topics that AI-generated articles tend to cover. For this illustrative analysis, we focus on pieces published around earnings announcements, a period that tends to attract more messaging activities on both platforms.[20] Panel A of Figure A1 reports

[19] A higher Fog Index need not imply poorer communication. In a professional finance setting, AI-assisted writing may employ more formal and technically precise exposition, with longer sentences and more specialized vocabulary, which mechanically raises Fog.

[20] There are 7,965 articles published within a three-day window following earnings announcements (days

the topic shares and associated word clouds for AI and human Seeking Alpha articles. AI articles are substantially more likely to load on the topic we label "Business Performance and Growth," which emphasizes forward-looking language about revenue growth, margin expansion, and management outlook. In contrast, human-authored articles more often concentrate on "Financial Results," a topic dominated by concrete reporting of metrics such as EPS, EBITDA, and quarterly figures. Together, these patterns suggest that AI-assisted content on Seeking Alpha tilts toward broader, forward-looking narrative framing, whereas human content places relatively more weight on granular earnings recap.

Panel B of Figure A1 presents the corresponding analysis for WallStreetBets. The topic mix differs sharply from Seeking Alpha, with both AI and human content frequently focused on trading activity and positions, as well as speculation related to AI and semiconductor-linked stocks. Relative to human content, pieces with AI content are more likely to mention financial results and accounting-style metrics, consistent with AI filling in a more fundamentals-oriented style that is less common in organically written WallStreetBets discussion.

4 AI Adoption in Content Production

In this section, we empirically investigate the determinants of AI adoption for the creation of financial content. We begin our analysis by focusing on Seeking Alpha, for which we obtain biographical information for a majority of authors in our sample. We perform author-level tests of how author backgrounds such as whether English is the primary language, education, and professional backgrounds are associated with AI usage. These tests are motivated by Hypothesis 1a, which predicts higher AI adoption when authors face higher information production costs or learning frictions, and when AI's pedagogical capabilities are particularly valuable for acquiring and interpreting information.[21]

[0, +3] relative to the announcement date, 9.5% of the SA article sample). A fraction of 91.2% of these articles reference earnings in some form. We apply non-negative matrix factorization (Lee and Seung, 1999) and fit separate topic models for each platform. We extract ten topics per corpus and assign each article to its dominant topic based on maximum loading.

[21] We do not conduct analogous author-level tests for WallStreetBets because author metadata on Reddit is substantially more limited and less reliably linkable to stable characteristics. In our WSB data, authors are

We then turn to article level analysis and compare the AI adoption patterns between Seeking Alpha and WallStreetBets and test whether observed AI deployment is associated with circumstances when information costs are high or during episodes of intense retail trading and speculative activity. These tests help discern whether observed AI deployment is more consistent with the *reason enhancement* channel (Hypothesis 1a) or with the *information distortion* channel (Hypothesis 1b).

4.1 Seeking Alpha Author Characteristics and AI Adoption

Motivated by Hypothesis 1a, this subsection examines whether AI adoption on Seeking Alpha is systematically related to author characteristics that proxy for language frictions, learning costs, and the intensity of analytical tasks.

We extract author characteristics from Seeking Alpha biography pages, which are available for 96.3% of the 1,998 authors in our sample. A substantial fraction of the authors hold professional certificates: CFA (Chartered Financial Analyst, 7.5%), FRM (Financial Risk Manager, 0.4%), CPA (Certified Public Accountant, 1.9%), CFP (Certified Financial Planner, 0.4%), and CAIA (Chartered Alternative Investment Analyst, 0.1%). Similarly, many of the authors are investment professionals, including portfolio managers (5.8%), traders (13.9%), financial advisors (11.1%), and financial analysts (8.8%). In terms of educational background, 3.8% hold a PhD and 4.8% hold an MBA degree. Additionally, 8.8% of authors on Seeking Alpha report holding a bachelor's degree (Internet Appendix A provides additional methodological details). We further identify authors likely to be native English speakers (15.2%) using name-based inference. We run the regression

$$AIProb_{ijt} = \beta_1 AuthorChars_j + Controls + FE + \epsilon_{ijt}, \tag{1}$$

where the dependent variable $AIProb_{ijt}$ is the AI probability score (in percentage points) of article j by covering stock i on day t. $AuthorChars_j$ are indicator variables that correspond to

typically identified only by a username (often anonymous or throwaway), and the platform does not provide biographical attributes comparable to Seeking Alpha profiles.

a given author characteristic. We control for article characteristics, including length, sentiment, complexity, fog index, graphical, and quantitative. We also include day and stock fixed effects to absorb time-varying and stock-level confounds.

Table 2 reports the results. In column 1, the coefficient on *English as Primary Language* is negative and significant, indicating that authors who are inferred to be non-native English speakers exhibit significantly higher AI usage, by 2 percentage points (pp), relative to native speakers. This is consistent with AI lowering language-related production costs and supporting comprehension and expression through its pedagogical functions. Column 2 reports that professional certification is also associated with lower AI probability scores: *Certificated* corresponds to a 1.78 pp decrease in AI probability. Similarly, authors who are *Portfolio Manager*, *Trader*, and *Advisor* are significantly less likely to use AI, by −3.69 pp, −0.714 pp, and −1.97 pp, respectively. This is consistent with their possession of specialized domain expertise that reduces the marginal cost of information acquisition. An exception is financial analysts, who are substantially more likely to use AI, by 4.475 pp. Similarly, PhD holders also exhibit higher AI adoption, by 7.767 pp. We also observe higher adoption among authors with MBA education, by 1.347 percentage points. These patterns suggest that AI complements analytical work for authors with more research-oriented backgrounds (analysts and PhDs) and for generalists (MBAs), by helping users learn about unfamiliar settings, organize complex information, and translate insights into coherent written form.

We obtain very similar results pooling all these attributes in the same regression in column 9, although the coefficient on *Trader* becomes insignificant. Overall, consistent with Hypothesis 1a, these patterns suggest that AI adoption on Seeking Alpha is more prevalent among authors facing language or learning frictions, and among authors for whom the pedagogical benefits of AI in acquiring, structuring, and synthesizing information are particularly valuable,

4.2 Information Costs and AI Adoption

In this subsection and the next, we turn to tests that more directly compare dynamic AI adoption patterns across Seeking Alpha and WallStreetBets. To isolate platform-level mechanisms from fixed author traits, we either control explicitly for author characteristics or include author fixed effects in the analysis.

We begin by testing whether AI adoption is higher when authors face higher information production costs, consistent with a pedagogical role of AI in helping authors acquire, interpret, and synthesize unfamiliar information. In particular, if AI lowers the cost of learning about a firm and organizing dispersed information into a coherent narrative, then authors will be more likely to use AI when covering unfamiliar firms and when the public information environment is sparse.

We estimate a modified equation (1), replacing the key explanatory variable with an indicator variable, $Unfamiliarity_{ijt}$, that equals one if the author/user covers the stock i for the first time in the past six months. *Controls* include stock size, book-to-market ratio and institutional ownership, as well as the author's past six-month activity (number of articles, number of AI articles, and average comments per article in the first 20 days after publication). In addition to stock and day fixed effects, we also add author fixed effects to control for time-invariant author characteristics. Standard errors are two-way clustered by stock and day.

The results in Panel A of Table 3 support Hypothesis 1a on the Seeking Alpha platform. On SA (column 1), first-time firm coverage is associated with a 2.32 percentage points (pp) higher AI probability ($t = 9.42$). This effect is economically significant relative to a mean of 10.65 pp on SA: when venturing into new domains, authors are 22% more likely to use AI. To address unobserved author characteristics (such as educational and professional background, writing style, technological familiarity, or risk preferences) that may influence adoption, column 2 includes author fixed effects. We find that the effect of first-time firm coverage remains positive and robust, though smaller in magnitude, indicating that within-author variation in informational familiarity continues to predict AI use.

We also find that firm size is negatively related to AI adoption: SA authors are more likely to use AI when writing about smaller firms, which typically suffer from limited information availability. For author characteristics, those with fewer articles in the past six months, those with prior experience using AI, and those whose articles generate more engagement (measured by comments per article) are all more likely to adopt AI. These associations persist even after controlling for author fixed effects. Taken together, these patterns reinforce the interpretation that AI serves a pedagogical and organizational function on SA, helping authors learn about new firms, structure complex information, and translate research into coherent narratives when traditional information acquisition is more costly, in line with Hypothesis 1a.

In contrast, on WSB (columns 3–4), there is no statistically significant increase in AI adoption when contributors cover firm for the first time. This divergence indicates that, unlike on SA, the adoption of AI by WSB users is not primarily associated with efforts to overcome information frictions. Firm size is weakly and negatively correlated with AI adoption, and book-to-market ratio and institutional ownership show no systematic relationship. For WSB author characteristics, column 3 shows higher AI adoption among users with fewer past articles, prior AI use, and lower comment engagement, but these associations weaken once author fixed effects are introduced. These results suggest that on WSB, AI adoption is not systematically tied to learning about new firms or filling gaps in traditional information sources. This tends to oppose the reason-enhancement hypothesis on this platform.

We further test Hypothesis 1a by examining whether AI adoption rises during periods when firm-specific information is more scarce. We measure information availability using the daily count of high-relevance news items from RavenPack and proxy ex-ante information supply by the number of analysts following the stocks. We estimate a modified equation (1), replacing the key explanatory variable with information scarcity measures: $News_{it}$ is the logarithm of one plus the number of highly relevant Dow Jones news articles (with a Ravenpack relevance score above 90 on firm i in the last three days), and *Analysts* is the logarithm of one plus the number of analysts following the stock i in the most recent quarter. Both

variables are standardized to zero mean and unit variance. Control variables, fixed effects, and standard error clustering are identical to the specification in Panel A of Table 3.

Results in Panel B of Table 3 are also consistent with Hypothesis 1a on SA. On the SA sample (columns 1-2), AI adoption is significantly less likely when news or analyst coverage is abundant. A one-standard-deviation increase in news coverage is associated with a 0.32 to 0.21 percentage point decrease in AI probability, while a one-standard-deviation increase in analyst following corresponds to a 0.73 to 0.33 percentage point decrease in AI probability. These findings indicate that SA contributors are more likely to use AI when traditional information sources are limited. In contrast, no such relationship is observed on WSB (columns 3-4), and in some specifications analyst coverage is even positively associated with AI adoption. This divergence highlights that information scarcity is not a primary driver of AI use on WSB.

Overall, these findings are consistent with the idea that AI flattens the learning curve in SA, lowering the marginal cost of research when authors lack deep prior knowledge. However, these patterns are absent on WSB, underscoring that AI adoption differs across platforms, with SA usage aligning with the reason-enhancement motive and WSB usage instead consistent with a belief distortion motive, to which we now turn.

4.3 Retail Buying Intensity and AI Adoption

We next test Hypothesis 1b, which posits that AI adoption on lightly moderated and speculative platforms such as WSB is associated with retail-driven speculation. Under this hypothesis, users turn to AI when retail interest is elevated, using it to craft persuasive or attention-grabbing narratives that resonate with peers and reinforce community dynamics. In this belief-distortion view, AI amplifies sentiment and narrative transmission around retail buying surges, rather than primarily serving to reduce information frictions.

We therefore test whether AI adoption is associated with intense retail investor buying.

We estimate the following article-level regression.

$$AIProb_{ijt} = \beta_1 IntenseRetailBuy_{i[t-3,t-1]} + Controls + FE + \epsilon_{ijt}, \quad (2)$$

where $IntenseRetailBuy_{i[t-3,t-1]}$ equals one if retail net buying during days [t-3, t-1] falls in the top 10% of the cross-sectional distribution. Following Boehmer et al. (2021), retail net buying is defined as the number of shares bought minus the number of shares sold by retail investors, scaled by total retail trades. A spike in retail net buying therefore serves as a proxy for strong retail buying pressure. Control variables, fixed effects, and standard error clustering are identical to the specification in Table 3.

The results, presented in Table 4, are consistent with Hypothesis 1b. For messages on WSB (columns 3–4), the coefficient on *IntenseRetailBuy* is positive and highly significant, both statistically and economically. A surge in retail buying predicts a 1.59 to 1.99 percentage point increase in a content's AI probability, which corresponds to a 44% to 55% rise relative to the platform's sample average. This indicates that WSB users are substantially more likely to adopt AI when a stock is experiencing intense retail buying.

In contrast, this relationship is not observed on Seeking Alpha. As shown in columns 1 and 2, the coefficient for *IntenseRetailBuy* is statistically insignificant for SA articles. This divergence suggests that while AI adoption on SA is more closely associated with information frictions and cost-saving (reason enhancement), AI adoption on WSB is more strongly linked to retail-driven surges and speculative dynamics (belief distortion).

5 AI Rhetoric and User Responses

We now turn to how AI-generated content shapes user responses through rhetoric. In our framework, AI can engineer distinct persuasive styles, which we summarize using Aristotle's three rhetorical modes. As defined earlier, we distinguish persuasion modes of *logos*, *ethos*, and *pathos*.

Under Hypothesis 2, AI flexibly engineers rhetorical content that aligns with platform

incentives: it generates high-*logos* and high-*ethos* content on curated platforms that reward analytical rigor, while it is produces high-*pathos* content on lightly regulated and more speculative platforms. Under the behavioral interpretation of Hypothesis 3, AI further amplifies belief distortion by strengthening sentiment contagion and the social transmission of biased narratives. We test these two hypotheses by first documenting cross-platform differences in the rhetorical composition of AI-generated content and then examining how user-comment sentiment responds to AI content on each platform.

5.1 Use of Different Rhetorical Modes

For each piece of content, we use GPT-5.1 to estimate the fraction of text attributable to each rhetorical mode, *logos*, *ethos*, and *pathos*, which we refer to as its rhetorical *Share*. We also use GPT-5.1 to estimate the effectiveness of each mode within an article and to construct an article-level measure of persuasive effectiveness.

The mode-specific shares are mutually exclusive and sum to one for each article. Mode effectiveness is measured on a 0 to 10 scale and we define a piece's overall persuasive effectiveness, *Effectiveness*, as the sum across modes of each mode's share multiplied by its effectiveness. Internet Appendix B describes classification details.[22]

Appendix Table A6 provides illustrative examples. Both AI-generated and human-author articles on Seeking Alpha emphasize *logos*, i.e., structured valuation arguments, and *ethos*, i.e., language that signals expertise and the credibility of the author or the author's information sources. In contrast, relative to Seeking Alpha, AI-generated messages on WallStreetBets have a greater tilt toward *pathos*, characterized by emotionally charged language and narrative amplification aimed at mobilizing retail sentiment.[23]

[22]The model is instructed to prioritize classification into *logos* when content simultaneously conveys factual analysis and author credibility, assigning statements grounded in data, valuation, or explicit reasoning to *logos* rather than *ethos*. The *ethos* category is reserved for explicit credibility signals, such as references to professional experience, past performance, or reputational claims, while *pathos* captures emotional language, narrative framing, and affective appeals.

[23]For example, an AI-generated Seeking Alpha article classified as 88% *logos* (effectiveness score 9/10) states: "Analysts predict a 13.5% decline in earnings to $0.46 per share compared to last year, while revenue is forecasted to increase by 7.1% to $124.7 billion." The rhetorical classifier notes that this message is "overwhelmingly

Figure 2 compares rhetorical shares (left axis) and article effectiveness (right axis) for AI-generated and human-authored content on each platform. On Seeking Alpha (Panel A), AI and human content exhibit nearly identical rhetorical profiles, with approximately 76–77% *logos*, 10% *pathos*, and 13–14% *ethos*, and identical effectiveness scores of 7.7. This similarity reflects that AI-generated content largely conforms to platform norms that emphasize analytical rigor and source credibility, rather than altering rhetorical composition.

The pattern on WallStreetBets differs along three dimensions. First, relative to AI generated content on Seeking Alpha, AI messages on WallStreetBets allocate substantially more text to *pathos* (34%), compared with 9.8% for AI articles on Seeking Alpha. This cross-platform comparison is consistent with Hypothesis 2, which predicts greater reliance on emotionally charged language in speculative environments.

Second, within WallStreetBets, AI-generated messages allocate a larger share to *logos* than human-authored messages (54% versus 42%), while human-authored messages devote a higher share to *pathos*. Third, AI-generated pieces achieve higher overall persuasive effectiveness than do human-authored pieces (6.2 versus 5.6), indicating that AI deploys rhetoric more effectively on WallStreetBets even though it does not increase emotional content relative to the human baseline.

These within-platform patterns highlight an important nuance for Hypothesis 2. While the hypothesis emphasizes differences in rhetorical emphasis across platforms, AI-driven belief distortion on WallStreetBets does not operate solely through a higher share of emotional language relative to human-authored messages. Instead, the rhetorical engineering facilitated by AI may operate through effectiveness rather than composition: conditional on a given rhetorical mix, AI can craft more persuasive arguments.

data-driven and analytical, with only light emotional language." In contrast, an AI-generated WallStreetBets message classified as 90% *pathos* (effectiveness score 9/10) reads: "Imagine selling RKLB at $8, then watching helplessly as it rockets to $22." The classifier identifies this as "an emotional FOMO narrative with vivid regret and urgency." In the third example, we report a Seeking Alpha article classified by GPTZero as likely human-authored, in which *ethos* is the dominant mode (40% ethos, effectiveness score 9/10): "Those who have tracked us over time have recognized the accuracy of not only our short-term calls but also our long-term ones." The classifier explains that this message "primarily catalogs past market calls and trade outcomes to showcase the author's track record."

To test this, we sort messages into terciles based on their *pathos* share and examine, within each tercile, differences in rhetorical effectiveness between AI-generated and human-authored content. Appendix Table A7 reports the effectiveness of different persuasion modes for a given article and finds that the rhetorical advantage of AI is context-dependent. In the high-pathos tercile, AI significantly outperforms human messages in *pathos* effectiveness (7.37 versus 7.08), but shows little advantage in *logos*. In comparison, in the low-pathos tercile, AI-generated messages outperform human-authored messages primarily through superior *logos* effectiveness (6.60 versus 5.64). The middle tercile displays a hybrid pattern, with AI outperforming on both *logos* (4.86 versus 3.98) and *pathos* (7.23 versus 6.57). Taken together, these results indicate that belief distortion on speculative platforms need not arise from greater emotionality per se. Instead, they can arise when AI makes speculative narratives more persuasive—whether through emotional resonance or through arguments that appear more coherent or analytically grounded.

5.2 Sentiment Contagion

The measured effectiveness of messages in the previous subsection is based on the assessment of an LLM (GPT5.1). We now examine user responses to AI-generated persuasion on Seeking Alpha versus WallStreetBests. Hypothesis 3 focuses on this downstream response. Persuasion through *logos* and *ethos* primarily operates by organizing information and reinforcing source credibility, and is therefore less likely to generate belief overshooting. In contrast, persuasion through *pathos* relies on emotional salience and narrative intensity, which can amplify behavioral biases. As a result, AI-generated content that emphasizes *pathos* is more likely to generate contagion of irrational sentiment, especially on speculative platforms where users are more susceptible to social influence.

As WallStreetBets users are generally less sophisticated and more exposed to social influence than Seeking Alpha users, we expect overshooting associated with AI use to be more common on WallStreetBets. To test this idea, We examine whether, on the two platforms, AI-generated content strengthens the alignment between the sentiment of an initial post and

the sentiment expressed in subsequent user comments.

Specifically, we test whether AI-generated content is associated with stronger sentiment contagion between articles and subsequent user comments using the following specification:

$$\begin{aligned} CommentSentiment_{ij[t,t+10]} =& \beta_1 AIProb_{ijt} \times Sentiment_{ijt} + \beta_2 Sentiment_{ijt} + \\ & \beta_3 AIProb_{ijt} + Controls + FE + \epsilon_{ijt} \end{aligned} \tag{3}$$

where $CommentSentiment_{ij[t,t+10]}$ is the average sentiment of comments following article j about stock i over the window $[t, t+10]$. We set *CommentSentiment* to zero when there are no comments. (Our results are robust to restricting the sample to articles and messages with comments; see Internet Appendix Table IA1.) $Sentiment_{ijt}$ is the sentiment of the article itself, and $AIProb_{ijt}$ is the AI probability score. The interaction term β_1 captures whether AI content creates stronger sentiment transmission from articles to comments. Control variables include article characteristics (length, complexity, Fog Index, and graphical and quantitative proportions) as well as the stock and author controls in Table 3. Fixed effects and standard error clustering are identical to the specification in Table 3.

Table 5 presents the results. For Seeking Alpha, we find no significant association between AI and sentiment contagion. Columns 1 and 2 show that the interaction coefficients are -0.022 and -0.046 and insignificant, so there is no evidence that AI-generated content is associated more strongly than human-written content with sentiment contagion. The baseline sensitivity of comment sentiment to article sentiment remains strong at 0.217, but AI does not appear to strengthen this relationship.

In contrast, on WallStreetBets, AI substantially amplifies sentiment transmission. In columns 3-4, the interactions coefficients are positive and statistically significant, at 0.090 and 0.135, respectively. Given the baseline sensitivity of 0.080 and 0.067, this represents a doubling of sentiment influence when content is AI-generated. Combined with our evidence that AI-generated WallStreetBets messages disproportionately emphasize *pathos* rel-

ative to those on SeekingAlpha, the results point to emotional persuasion, rather than *logos*- or *ethos*-based persuasion, as the more plausible channel behind the stronger sentiment contagion.Consistent with Hypothesis 3, these findings suggest that AI-assisted rhetorical engineering may amplify social transmission bias (as modeled in Hirshleifer (2020) and Han, Hirshleifer, and Walden (2022)) on speculative platforms. However, the tighter link between post sentiment and comment sentiment could also reflect rational updating, wherein users more closely process and respond to the information conveyed in the initial post.

To distinguish rational updating from amplification or the spread of simplistic ideas, or affective amplification, wherein emotionally charged narratives induce belief overshooting and sentiment contagion, we test whether sentiment alignment associated with AI posts is linked to price distortions or speculative return patterns. Alignment should only contribute to price bubbles or lottery-like return outcomes if alignment derives from overreaction to sentiment of the original post. Motivated by this distinction, we next turn to tests of asset pricing patterns and market outcomes linked to AI-generated content.

6 AI-Generated Content and Capital Market Outcomes

In this section, we investigate the association between AI-generated content and capital market outcomes. We examine whether its adoption is associated with improved market quality and price discovery (Hypothesis 4, consistent with information enhancement) or with increased speculative trading and lottery-like return patterns (Hypothesis 5, consistent with belief distortion). We first conduct a direct test of informativeness by examining the return predictive power of AI article sentiment. Next, we consider the association between AI content and broad indicators of market quality: abnormal trading volume, volatility, and bid-ask spreads. We then assess whether AI enhances the informativeness of retail order flows. Finally, we test for evidence of speculative price dynamics by investigating whether AI content predicts “lottery-like” return events.

6.1 Informativeness of AI Articles

We proxy for an article's fundamental informativeness by examining the extent to which its sentiment aligns with subsequent stock returns. Under the reason-enhancement hypothesis, if an article conveys fundamentally accurate signals, and if this information is not fully and immediately incorporated into market price, then we expect a positive association between article sentiment and future returns. That is, when an article expresses bullish sentiment, future returns should, on average, be positive, and vice versa. By the same token, if bullish sentiment is followed by neutral or negative returns, the article is less likely to contain fundamentally accurate information.

Again, this prediction is premised on such information not being fully and immediately incorporated into market prices. This is consistent with limited attention and trading frictions, and with a large empirical literature documenting underreaction to earnings announcements and analyst forecast revisions (e.g., Bernard and Thomas 1989; Hirshleifer, Lim, and Teoh 2009; Stickel 1991; Chan, Jegadeesh, and Lakonishok 1996).[24]

To test these implications, we estimate the following article-level regression:

$$\text{CAR}_{i[t+1,t+5]} = \beta_1 AIProb_{ijt} \times Sentiment_{ijt} + \beta_2 Sentiment_{ijt} + \beta_3 AIProb_{ijt} + Controls + FE + \epsilon_{ijt}, \tag{4}$$

where $\text{CAR}_{i[t+1,t+5]}$ is the cumulative market-adjusted abnormal return over the subsequent five trading days, $AIProb_{ijt}$ is the AI probability score assigned to article j about firm i on day t, and $Sentiment_{ijt}$ is the article's sentiment score. The coefficient on *Sentiment* captures the baseline effect—the extent to which the sentiment of human-authored articles align with future returns. The key variable of interest, the coefficient β_1 on the interaction term, which measures the incremental tendency of AI-generated content relative to human-generated articles to predict higher future returns. *Controls* include the article characteristics (length,

[24] An article whose bullish sentiment systematically predicts negative subsequent returns can still be informative for sophisticated traders who profit by trading against its views. Our definition of informativeness here is narrower: we focus on whether the article contains fundamentally accurate information.

sentiment, complexity, Fog Index, and graphical and quantitative proportions). Additionally, we also include stock characteristics: size, book-to-market ratio, institutional ownership, analyst coverage, past news, an earnings announcement dummy, and past one-month return and volatility. *FE* corresponds to firm and day fixed effects.

Table 6 presents the results. Columns 1–2 show results for SA articles and columns 3–4 for WSB messages. In column 1, the coefficient on *Sentiment* is −0.149 (significant at the 1% level), indicating that the sentiment of human-written SA articles *negatively* predicts returns in our sample. The coefficient remains negative and significant even if *AIProb*-related variables are excluded from the regression. This finding, estimated with our December 2022 to December 2024 sample, differs from work with earlier samples in which SA articles were found to be informative (Chen et al. 2014), but is consistent with recent work documenting the decay of the value of social media signals following the 2021 GameStop event (Bradley et al. 2023).

Crucially, the interaction term, $AIProb \times Sentiment$, is positive and statistically significant in predicting the next five-day returns. Economically, for a one-standard-deviation increase in (bullish) sentiment, a fully AI-generated article ($AIProb = 1$) predicts 35 basis points higher one-week-ahead returns than does a human-generated article ($AIProb = 0$). This result is consistent with the sentiment of AI articles containing useful information about future fundamental prospects.

We next test whether these return patterns simply reflect temporary sentiment-driven price pressure. If AI article sentiment is associated with short-lived mispricing, its initial return predictability should reverse as the distortion corrects. Column 2 finds no reversal in the subsequent $[t+6, t+10]$ window—the interaction coefficient is insignificant—indicating that AI article sentiment is not followed by correction. This reinforces the interpretation that AI-generated content on SA conveys value-relevant information rather than noise, consistent with the reason-enhancement hypothesis.

For r/WallStreetBets messages (columns 3-4), however, the interaction coefficients are insignificant across both time windows. This suggests that AI content on this platform is not

informative about future returns.[25] This divergence between the findings for SA and WSB aligns with our earlier findings that the influence of AI on returns differs across platforms.

6.2 Informativeness of Retail Orders

The predictive power of AI articles for returns on SA suggests that such articles contain value-relevant information that is not immediately incorporated into prices. This raises the question of whether such signals help attentive investors—particularly rational retail traders facing high information costs—earn higher returns.

As shown by Kaniel, Saar, and Titman (2008) and Boehmer et al. (2021), retail net buying is associated with positive subsequent abnormal returns, which indicates that net retail buying is informative or is associated with liquidity provision. Building on this evidence, we examine whether the presence of AI-generated content strengthens the return predictive power of retail net buying.

If AI adoption reflects rational information acquisition (Hypothesis 4), retail net buying should be more predictive of future returns on days when AI-generated content is present. Alternatively, under the belief-distortion channel, AI deployment should not enhance—and may even weaken—the link between retail order imbalance and subsequent performance.

To test these predictions, we next turn to stock-level analysis and run the following stock-day panel regression to predict returns:

$$\begin{aligned} \mathrm{CAR}_{i[t+1,t+5]} = {} & \beta_1 AIDay_{it} \times \mathrm{OIB}_{it} + \beta_2 HumanDay_{it} \times \mathrm{OIB}_{it} \\ & + \beta_3 AIDay_{it} \times Sentiment_{it} + \beta_4 HumanDay_{it} \times Sentiment_{it} \\ & + \beta_5 AIDay_{it} + \beta_6 HumanDay_{it} + \beta_7 \mathrm{OIB} + \beta_8 Sentiment_{it} \\ & + \beta_9 Attention_{it} + Controls + FE + \epsilon_{it}, \end{aligned} \tag{5}$$

where $\mathrm{CAR}_{i[t+1,t+5]}$ is the market-adjusted cumulative abnormal return from $t+1$ to $t+5$.

[25] The higher R-squared for the WSB sample than the SA sample likely reflects differences in firm coverage: as shown in Table 1, SA incentivizes coverage of smaller, under-followed stocks, whereas WSB messages tend to focus on larger firms.

The main variable of interest is the interaction between $AIDay_{it}$, an indicator for at least one AI article for stock i on day t, and OIB_{it}, the retail order imbalance calculated as the ratio between the net retail buys and total retail volume (in shares) following Boehmer et al. (2021). The coefficient on this term (β_1) captures whether retail order flows are more informative on AI days.

We also include $HumanDay_{it}$, an indicator for firm-days with at least one human-generated article, where human articles are identified by GPTZero as having the highest probability of being "HUMAN-ONLY" content. We further control for the total number of articles ($Attention_{it}$) as a proxy for social media attention and average sentiment of all content for that stock on that day (*Sentiment*). *Controls* is a vector that includes size, book-to-market ratio, institutional ownership, analyst coverage, past news, an earnings announcement dummy, and past one-month return and volatility. We further include platform-specific interactions between $AIDay_{it}$ and platform-level sentiment. Standard errors are two-way clustered by stock and day.

We estimate this regression separately for the SA and WSB samples and then for a pooled sample that includes all stocks that appear on either platform.[26] Table 7 presents the results. Across all specifications, we find a positive and statistically significant coefficient on the $AIDay^{SA} \times \text{OIB}$ interaction in the SA sample. The coefficient of 0.431 in column 2 indicates that on days with AI content on Seeking Alpha, a one-standard-deviation greater OIB predicts an additional 43.1 basis points in cumulative abnormal returns over the subsequent five trading days, compared to days without AI content. This association is economically substantial and statistically significant. In contrast, the interaction between $HumanDay^{SA}$ and OIB is much smaller (0.053) and statistically insignificant, indicating that human articles do not enhance the return predictive power of retail order flows. These findings are consistent with Hypothesis 4: AI content on SA appears to enhance the informativeness of

[26]The SA sample includes all stocks covered by at least one Seeking Alpha article during our sample period, and the WSB sample includes all stocks mentioned in at least one WallStreetBets message. For the pooled analysis, on days when a stock has no content on a given platform, we set *AIDay*, *Sentiment*, and *Attention* to zero for that platform, following Cookson et al. (2024).

retail trading, in line with the reason-enhancement channel.[27]

In contrast, the interaction coefficient for the WSB sample ($AIDay^{WSB} \times$ OIB) is positive but not statistically significant (columns 3–4), suggesting no detectable association between AI presence and the strength of the OIB-return relationship on WSB. The $HumanDay^{WSB} \times$ OIB interaction is similarly insignificant (around 0.070), indicating that neither AI nor human content on WSB enhances the informativeness of retail order flows. This difference between platforms may reflect the distinct nature of their user bases and content types, with the more analytical focus of Seeking Alpha being associated with improved price discovery from AI-assisted information synthesis. The combined model (columns 5–6) shows the same pattern. The SA interaction remains significant while the WSB interaction remains insignificant.

Overall, these results suggest that AI deployment on SA is consistent with Hypothesis 4: AI content is associated with more informative retail order flows, consistent with AI-generated content helping rational and attentive retail investors on this platform better synthesize and act upon complex information. The contrast with human content—for which the interaction effect is much smaller and insignificant—indicates that this finding reflects AI-specific mechanisms rather than general content effects. On WSB, neither AI nor human content appears to enhance the informativeness of retail trading.

6.3 Volume, Volatility, and Bid-ask Spreads

We next examine how the presence of AI-generated content is associated with market quality measures such as trading behavior and stock liquidity. We estimate the following stock-day panel regression using all trading days:

$$MarketQuality_{i[t+1,t+5]} = \beta_1 AIDay_{it} + \beta_2 HumanDay_{it} + \beta_3 Attention_{it} + \beta_4 Sentiment_{it} + Controls + FE + \epsilon_{it}. \quad (6)$$

[27] Table 7 also indicates that for our December 2022 to December 2024 sample the unconditional OIB-return relation is weak (though positive in sign). Our main result concerns the incremental effect: OIB is significantly more informative on SA AI days relative to other days.

MarketQuality refers to the following trading activity variables measured over the subsequent five days: abnormal trading volume, return volatility, and the bid-ask spread. All dependent variables are standardized, so the coefficients represent changes in standard deviations. The key independent variable is $AIDay_{it}$, an indicator for firm-days with at least one AI-generated article for a given platform. Since bid-ask spreads are highly persistent (Chordia, Sarkar, and Subrahmanyam 2005), we also control for the average daily bid-ask spread in the past one month when *MarketQuality* is measured by bid-ask spread. For the other explanatory variables and fixed effects included, definitions follow those provided in Table 7.

The findings, presented in Table 8, indicate important differences in the relation of AI to trading and capital market outcomes across platforms. Columns 1–3 examine Seeking Alpha: over the next five trading days, AI presence is associated with a 0.018 standard deviation (SD) decrease in effective bid-ask spreads (or a 5.3% decrease relative to the sample mean). This pattern of tighter spreads is consistent with the reason-enhancement channel and Hypothesis 4: AI-generated content on SA is associated with improved price discovery, which reduces uncertainty and adverse selection costs borne by liquidity providers. The coefficients on abnormal volume and volatility are also negative although insignificant.

On WSB, shown in columns 4–6, we observe a markedly different pattern for AI content. AI presence is associated with a 0.086 SD increase in abnormal volume (or a 161.0% increase relative to the near-zero sample mean) and a 0.074 SD increase in volatility (or a 10.5% increase relative to the sample mean).[28] The divergent relationship between AI content and market outcomes on the SA and WSB platforms underscores the critical role of the institutional and behavioral environment in shaping AI's market influence.

In the combined model (columns 7–9), where we pool observations from both platforms, the opposing effects across platforms persist. For SA, AI content is associated with a 0.016 SD decrease in spreads (or an 4.7% decrease relative to the sample mean), while the

[28]The coefficient on $HumanDay^{WSB}$ for volatility is −0.066 (significant at the 1% level). The negative coefficient represents the effect of human content conditional on total attention. In an alternative specification omitting the $Attention_{it}^{WSB}$ control, $HumanDay^{WSB}$ becomes positive and significant, indicating that the unconditional association of human-generated pieces with volatility is positive, similar to AI content.

coefficients on volume (0.015 SD) and volatility (0.022 SD) are positive but not statistically significant.

Overall, the across-platform comparison suggests that the relationship between AI with market variables depend on institutional context. On SA, where editorial review ensures content quality and with relatively sophisticated user base, AI is associated with value-relevant content and greater liquidity, consistent with Hypothesis 4.[29] On WSB, where anyone can instantly post AI-generated narratives, the same technology is linked to higher volatility and abnormal volume in subsequent periods, consistent with Hypothesis 5 and the belief-distortion channel.

6.4 Lottery-like Return Events

We have seen that the combination of elevated trading volume and volatility following AI-generated content on WSB is consistent with speculative trading. An alternative explanation is that these patterns are instead driven by large fundamental information shocks. To differentiate these possibilities, we next test whether AI content predicts extreme, lottery-like return events—sharp price run-ups. Such patterns of sharp price run-up followed by return reversals have been identified by past studies as characteristic of overvaluation (e.g., Kumar, Page, and Spalt 2011; Bali, Cakici, and Whitelaw 2011).

Following Bali, Cakici, and Whitelaw (2011) and Bali et al. (2025), we identify two types of events that capture lottery-like payoffs. A MAX event occurs when a stock's daily return is the highest over a trailing 21-trading-day window. A lottery event represents an even more extreme outcome: it is a MAX event and the MAXRET falls into the top decile of its daily

[29] This interpretation of AI's beneficial role on SA is broadly consistent with the findings of Bradshaw et al. (2025), who document that the introduction of ChatGPT is followed by a better overall information environment for stocks covered on SA, as reflected, for example, in lower spreads and higher liquidity. However, our interpretation differs in the comparison between AI-generated with human-generated content. Bradshaw et al. (2025) argue that the lower abnormal trading volume and volatility levels around the publication of AI articles, compared to human articles, indicates that AI articles are *less* informative. We argue that if the stock market underreacts, then informativeness is also measured by the extent to which article sentiment positively predicts subsequent returns. This allows for the possibility that the market does not immediately and fully incorporate that information into prices.

cross-sectional distribution, where MAXRET is the 21-day highest return. These events reflect the lottery-like characteristics that attract sentiment-driven retail traders but are less likely to be associated with fundamental information arrival. We estimate the following stock-day logistic regressions to predict the probability of these events:

$$\begin{aligned} \text{MAX/}Lottery_{i[t+1,t+5]} = {} & \beta_1 AIDay_{it} + \beta_2 HumanDay_{it} + \beta_3 Attention_{it} + \beta_4 Sentiment_{it} \\ & + Controls + FE + \epsilon_{it}, \end{aligned} \tag{7}$$

where the dependent variables are indicators for whether stock i experiences a MAX or lottery event either during the subsequent five-day window $[t+1, t+5]$. The explanatory variables, fixed effects, and standard error clustering are identical to the specification in Table 8.

Table 9 presents striking results. On WSB, AI content predicts the occurrence of both types of extreme events. For MAX events, AI presence is associated with an increase in odds by 18.0% ($e^{0.166} - 1$, column 3) in the next five days. The lottery event results are even more dramatic: AI content is associated with an increase in odds by 44.2% ($e^{0.366} - 1$, column 4) for the next five days. In the combined sample (columns 5–6), the WSB AI Day coefficients remain positive and significant, with odds increases of 18.6% ($e^{0.171} - 1$) for MAX events and 49.7% ($e^{0.403} - 1$) for lottery events. In contrast, $HumanDay^{WSB}$ shows no significant association with either type of extreme event, suggesting that the lottery-like dynamics are specific to AI-generated content rather than reflecting general messageing activity on WSB.

On SA, we find the opposite pattern. AI presence is associated with a *reduction* in lottery events: the coefficient of −0.336 (marginally significant at the 10% level, column 2) implies that AI content is associated with a 28.5% decrease in the odds of lottery events ($1 - e^{-0.336}$). Human-generated content shows a stronger negative association: $HumanDay^{SA}$ is associated with a 30.6% decrease in lottery event odds ($1 - e^{-0.365}$, significant at the 5% level). The coefficients for MAX events are also negative for both AI Day (−0.027) and Human Day (−0.065) on SA, though not statistically significant. In the combined model (columns 5–6), SA content

becomes statistically insignificant.

Extreme right-tail events are arguably symptomatic of price bubbles, i.e., speculative overvaluation, rather than of fundamental news. So the finding that AI presence predicts these return episodes on WSB supports the relevance of Hypothesis 5 and the belief-distortion channel on that platform. Conversely, the absence—and in some cases reduction—of lottery-like events on SA is consistent with Hypothesis 4 and the reason-enhancement channel.

7 Further Analysis and Robustness

This section first characterize within-article human-AI interaction patterns, then exploit the introduction of no-AI enforcement on Seeking Alpha as a quasi-natural experiment to study how these interactions change. We also examine user engagement and disagreement following AI content and presents robustness tests.

7.1 Patterns of Human-AI Interaction

Our analysis thus far has focused on an article-level AI-probability measure, which does not reveal *how* AI is used within an article. In practice, AI can play multiple roles, including idea generation, background research, narrative framing, elaboration or polishing of human-written text, or near-complete substitution for human effort. Although we do not observe the production workflows of an author directly, the placement of AI-generated text within an article (e.g., the beginning, the body, the end, or throughout) provides indirect but informative evidence on how human and AI inputs are combined.

Different modes of AI use naturally imply different integration points within the text. For example, AI employed for broad drafting or rewriting should appear throughout an article, whereas AI used more selectively is likely to concentrate in particular sections. Motivated by this intuition, we use sentence-level AI probability scores from GPTZero to classify articles based on the positional distribution and clustering of AI-like sentences. This approach allows us to distinguish between episodic, assistive uses of AI and more pervasive forms of

substitution.

We define four AI usage categories (Internet Appendix C provides methodological details). First, *Uniform AI* captures articles in which AI probability is high and relatively flat across the full text, with little systematic variation by section; this pattern is consistent with pervasive AI involvement, either because the author relies on AI for most content generation or because AI is used to rewrite an initial human sketch or draft across the entire article.

Second, *Body-Loaded AI* describes articles in which AI probability is comparatively low near the beginning and end. This pattern suggests that authors selectively deploy AI for substantive discussion, such as background information, quantitative elaboration, or supporting analysis, while retaining greater human involvement in framing. Such usage is consistent with AI acting as a complement to human analysis by expanding or clarifying informational content.

Third, *Edge-Loaded AI* applies when AI probability is elevated at the start and/or end of the article, indicating that AI-generated sentences are concentrated in framing sections, such as crafting introductions or conclusions, while relying more heavily on human input for core analytical content.

Finally, *Interspersed AI* captures articles in which AI-like sentences are alternate frequently with human-like sentences throughout the text, with no strong concentration in any single section. This pattern is consistent with iterative human-AI revision at a granular, perhaps sentence-by-sentence level.

Appendix Figure A2 illustrates these four patterns using articles from Seeking Alpha.[30]

[30]Panel A corresponds to *Uniform AI* and features an article about Mastercard, with AI probability equal to 100% and essentially constant across all 61 sentences. Panel B corresponds to an article about Shake Shack and illustrates the *Body-Loaded AI* pattern: the author opens with a personal investment thesis and market context (AI probability 0.29 in the introduction), while AI-like sentences are concentrated in the analytical body, where the mean AI probability rises to 0.80 and covers topics such as licensing expansion, international growth strategy, and SG&A leverage. The conclusion returns to human-written assessment (AI probability 0.08). Panel C is an article about Coca-Cola and provides an example of *Edge-Loaded AI*: AI probability is high at the start (mean 0.80 in the introduction) and end (0.92 in the conclusion), but drops substantially in the body (with near zero AI probability for the first half), where the author presents the bulk of the original analysis. Finally, the UiPath article in Panel D exhibits the *Interspersed AI* pattern, with frequent alternation between AI-like and human sentences throughout. These four articles can be found at https://seekingalpha.com/article/4631187-mastercard-displays-technical-strength-solid-fundamentals, https://seekingalpha.com/article/4602352-shake-shack-shak-expansion-in-metro-overseas, https://seekingalpha.com/article/47019

We next present the distribution of AI usage categories for AI articles on each platform in Figure 3, based on 3,820 Seeking Alpha articles and 1,321 WallStreetBets articles (at least five sentences required) that are classified as “AI-ONLY" or “MIXED" by GPTZero. These distributions reveal pronounced cross-platform differences that complement our earlier article-level results. On Seeking Alpha, *Uniform AI* accounts for 62.9% of AI articles. At the same time, 37.1% of articles fall into non-uniform patterns that indicate more selective deployment of AI. Specifically, *Body-Loaded AI* accounts for 21.7% of articles and *Edge-Loaded AI* for 11.0%, with *Interspersed AI* representing 4.4%. Taken together, these patterns suggest that at a substantial fraction of Seeking Alpha authors use AI selectively, concentrating AI assistance in specific parts of the article rather than relying on it uniformly.

In contrast, WallStreetBets exhibits a more concentrated distribution. 84.7% of AI pieces fall into the *Uniform AI* category, while non-uniform patterns are relatively rare. *Body-Loaded AI* accounts for 7.8% of pieces, *Edge-Loaded AI* for 1.4%, and *Interspersed AI* for 6.1%. Thus, when WallStreetBets users employ AI, they are substantially more likely to rely on AI uniformly across the piece rather than deploying it selectively within the content.

These differences in AI usage patterns help clarify the mechanisms highlighted earlier. On Seeking Alpha, where authors face stronger reputational and economic incentives to provide informative analysis, AI is more often deployed in ways that complement human effort, such as expanding substantive discussion or refining narrative structure. On WallStreetBets, in contrast, AI use is dominated by uniform patterns that are more consistent with wholesale substitution, a mode of use that is less likely to enhance informational content and more likely to contribute to noise. These within-article patterns therefore provide direct evidence that institutional incentives shape not only whether AI is used, but also how it is used, with important implications for reason enhancement versus belief distortion in financial discourse.[31]

39-the-dividend-dynamo-how-coca-cola-delivers-both-growth-and-your-next-paycheck, and https://seekingalpha.com/article/4663360-uipath-promising-play-in-the-era-of-generative-ai.

[31]Our classification has limitations. The classification of patterns may not capture the full range of possible human-AI collaboration strategies. Moreover, our inferences from these patterns to authors’ production workflows is indirect. Nevertheless, the sharp and systematic cross-platform differences we document are unlikely to be driven solely by classification noise. The findings align closely with our broader findings on the associa-

7.2 A Quasi-Natural Experiment: Seeking Alpha No-AI Policy

AI use on Seeking Alpha declined sharply after the platform begins to enforce its no-AI policy in late October 2023, including the deployment of proprietary detection tools and the subsequent sanctions such as suspensions or removals of violating authors (see Figure 1). The share of AI-generated content fell beginning in December 2023 and stabilized at approximately 3-4% by March 2024. This policy change provides a quasi-natural experiment to examine how authors adjusted their use of AI in response to the policy change.

Figure A3 presents the distribution of AI usage categories before and after the ban, expressed as a share of AI articles (defined as AI-ONLY and MIXED articles by GPTZero). The compositional shift is substantial. The share of Uniform AI, articles exhibiting high AI probability throughout, declines from 68.6% to 55.1%, a reduction of 13.5 percentage points. In contrast, more selective usage patterns expand. Body-Loaded AI rises from 18.7% to 25.8%, and Edge-Loaded AI nearly doubles from 8.3% to 14.7%. Interspersed AI remain stable at roughly 4.5%.

These patterns suggest that enforcement disproportionately deterred or detected wholesale AI substitution, while more selective forms of AI assistance, in which AI supports specific sections rather than generating entire articles, either escaped detection or were tolerated. Among authors who continued to use AI after the ban, usage shifted toward more complementary and targeted deployment. If Uniform AI reflects lower-effort content substitution, whereas Body-Loaded and Edge-Loaded AI capture more deliberate human-AI collaboration, enforcement may have raised the average quality of AI-assisted content on the platform.

To examine this possibility, we re-estimate our key specifications by interacting the main variables of interest with the pre- and post-ban period indicators. When the main variables of interest are themselves interactions, such as $AIProb \times Sentiment$ or $AIDay \times$ OIB, we also include the corresponding lower-order interaction terms. Because the sample is split into pre- and post-ban periods, these subsample tests have weaker statistical power than the full-

tions between institutional context, the adoption of AI-generated content, and financial market outcomes.

sample analysis.

Appendix Table A8 reports the results. Several findings emerge. First, the positive association between AI adoption and author unfamiliarity with the covered firm (column 1) is concentrated in the pre-ban period, consistent with AI facilitating coverage expansion primarily before enforcement tightened. Second, AI adoption remains negatively associated with news and analyst coverage in both periods (column 2), although the association with news coverage weakens post-ban. This indicates that AI continues to fill information gaps in less-covered stocks even after enforcement. Columns 3 and 4 show no significant association between AI content and intense retail buying or sentiment alignment with subsequent comments in either period, mirroring the full-sample results.

Column 5 examines whether article sentiment predicts future returns more accurately when content is AI-assisted. The interaction between AI probability and article sentiment is insignificant in the pre-ban period but becomes positive and statistically significant post-ban. This pattern suggests that AI-assisted articles became more informative after the ban. A natural interpretation is that, by discouraging wholesale AI ghostwriting and inducing more selective use, the policy shifted AI usage toward more sophisticated applications that enhance, rather than replace, author analysis. Column 6 analyzes whether AI content is associated with more informative retail order flow. The positive coefficients indicate that AI days on Seeking Alpha are associated with stronger return predictability from net retail buying, although the estimates are not statistically significant, with t-statistics of 1.34 and 1.53.

Columns 7 and 8 show that AI content was not significantly related to abnormal trading volume or volatility in either period, consistent with the full-sample findings. Column 9 shows that AI days were associated with narrower bid-ask spreads both before and after the ban, indicating that the liquidity improvements linked to AI content persist under enforcement. Finally, Columns 10 and 11 examine whether AI content predicts large one-day price run-ups and lottery-like return outcomes. AI days were negatively associated with lottery-like returns in both periods, with marginal statistical significance only in the post-ban era.

Taken together, these results indicate that, as authors adapted to enforcement, those

who continued to use AI appeared to do so in more selective and sophisticated ways, producing content that was more informative and less speculative. The improvement in informativeness post-ban provides suggestive evidence that policies discouraging low-quality AI substitution can enhance the informational value of AI-assisted financial analysis.

7.3 AI Content, User Engagement, and Disagreement

We have seen that AI content on WSB is associated with greater sentiment alignment between the original posts and followup comments, whereas the pattern is not present for SA AI content. In this subsection, we examine whether AI-generated content is associated with user engagement and disagreement.

We estimate the regression

$$Engagement_{ij[t,t+10]} = \beta_1 AIProb_{ijt} + Controls + FE + \epsilon_{ijt}, \tag{8}$$

where $Engagement_{ij[t,t+10]}$ captures user responses to content j about stock i over the window $[t, t+10]$. We consider two measures: (i) $\log(1 + Comments)$, the natural logarithm of one plus the number of follow-up comments received, and (ii) *Disagreement*, the standard deviation of the follow-up comment sentiment scores, standardized to zero mean and unit variance. $AIProb_{ijt}$ is the AI probability of the initial article. Control variables, fixed effects, and standard error clustering are identical to the specification in Table 5.

Appendix Table A9 presents the results. On Seeking Alpha, AI-generated content is associated with lower follow-up comment volume (column 1: $\beta = -0.065$, $t = -2.48$), consistent with Bradshaw et al. (2025). Furthermore, AI content is also associated with lower follow-up comment disagreement (column 3: $\beta = -0.066$, $t = -1.99$). Economically, this corresponds to a 0.066 standard-deviation (or 4% of the mean) smaller disagreement, relative to human articles. This finding is consistent with the interpretation that AI-assisted articles help structure and clarify fundamental information, thereby reducing the scope for divergent interpretations among readers. In contrast, on WallStreetBets (columns 2 and 4), the coefficient on

AIProb is statistically insignificant for both engagement measures.

7.4 Robustness

We perform two sets of robustness checks. First, we verify that our estimated associations are not driven by confounding earnings announcements, or by the anticipation of such announcements. Second, we provide evidence that our WallStreetBets results are not sensitive to alternative sample construction choices.

Earnings Announcements A potential concern is that our main findings on market outcomes could be driven by the mechanical correlation between article publication and earnings announcement events themselves. If AI articles cluster around earnings dates more than human articles, the association between AI content and market outcomes could reflect the information shock from the announcement.

To address this possibility, we exclude all stock-day observations that fall within the $[-1, +1]$ window around an earnings announcement. This removes the period when announcement effects are most concentrated. Appendix Table A10 Panel A reports the results. In Panel B, we further assess whether our findings reflect anticipatory trading ahead of scheduled information releases by adding an indicator variable, $I^{Earnings}_{[t+1,t+5]}$, that equals one if a stock has an earnings announcement within the subsequent five trading days. The coefficient on this indicator is highly significant for volume, volatility, and lottery returns, confirming increased market activities in advance of earnings releases. The key finding is that the coefficients on $AIDay^{SA}$ and $AIDay^{WSB}$ remain similar to our main results in both panels. Together, these tests indicate that the relationship between AI-generated content and market outcomes does not derive solely from earnings announcements.

Alternative WSB Samples Our main analysis excludes WSB messages shorter than 50 words to ensure sufficient text length for reliable AI detection by GPTZero. To assess the sensitivity of our findings to this restriction, we rerun the full set of specifications on the complete

WSB message sample, assigning an AI probability of zero to observations below the 50-word threshold and including an indicator for messages exceeding 50 words in all regressions. The replication results are reported in Internet Appendix Table IA2, Panel A. Our main results remain robust, except for column 2, which indicates that AI adoption is higher for stocks that receive more analyst coverage, whereas in the original sample, the relationship is insignificant.

In Panel B, we assess the sensitivity of our findings to the exclusion of messages that are later deleted. While Seeking Alpha articles are written by identifiable contributors and published after an editorial review, WSB discussions are generated by users in real time, and a substantial number of messages are later removed by moderators or deleted by users.[32] We replicate the key analysis excluding deleted messages (although doing so introduces message-selection bias). Our main results remain robust.

8 Conclusion

This paper provides some of the first evidence on how generative AI is reshaping financial discourse across online platforms, and the association of generative AI use with investor behavior and market outcomes. Motivated by the idea that AI can either enhance analytical thinking or can distort beliefs, we compare two prominent platforms with sharply contrasting institutional environments: the curated Seeking Alpha, where contributors are primarily semi-professional analysts, and the largely unmoderated Reddit's r/WallStreetBets, dominated by retail traders with limited sophistication. This comparison allows us to explore how these institutional environments relate to cross-platform differences in the role of AI in the information ecosystem.

Using a state-of-the-art and validated AI detection tool, we find that AI adoption rose

[32]Among the 64,848 WSB messages, 4,417 (6.8%) are removed and 417 of those are detected as AI-generated by GPTZero. 47.3% of the removal is by platform moderators, and the rest is removed by the author. Common reasons for removal include spam, abuse, or rule violations. Removed AI-generated messages receive an average of 7.15 comments (median 3) in the five-day window after posting, whereas removed human messages receive 5.75 comments on average. For messages that are not removed, the corresponding averages are 9.40 (median 2) for AI messages and 3.72 (median 1) for human messages.

following the launch of ChatGPT, but with sharply divergent patterns. On Seeking Alpha, adoption aligns with a reasoning-enhancement channel. AI use is higher among non-native English authors, consistent with AI playing a pedagogical role in addressing language frictions and improving exposition. AI use is also higher among authors with research-oriented or generalist backgrounds (for example, analysts and advanced degree holders such as PhDs and MBAs), suggesting that AI is used to deepen, organize, and communicate complex analysis. In contrast, AI use is lower among domain specialists, consistent with smaller marginal benefits when expertise and specialized frameworks are already in place.

After accounting for author characteristics, we find that authors on Seeking Alpha rely on AI more when covering unfamiliar firms or firms with sparse public information. In contrast, on WallStreetBets, AI usage spikes during episodes of intense retail buying, consistent with the belief distortion channel in which AI facilitates emotionally resonant, persuasive, or bias-amplifying narratives.

We further find that AI is associated with different forms of persuasion on different platforms. On Seeking Alpha, AI-generated content conforms closely to existing norms of analytical rigor and source credibility, relying primarily on logical argumentation and authority cues, and achieving persuasion effectiveness comparable to human-authored articles. On WallStreetBets, in contrast, AI-generated pieces are more emotionally charged than AI pieces on SeekingAlpha. Within that platform, AI pieces are more persuasive than human pieces on average, and AI pieces are linked to stronger sentiment contagion. These findings are consistent with AI augmenting the salience and transmissibility of speculative or affectively charged narratives on this platform.

The association between AI pieces and market outcomes also differs substantially across the two platforms. On Seeking Alpha, AI article sentiment positively predicts future returns, and following AI deployment, bid-ask spreads narrow and retail order flows become more informative. This evidence is consistent with improved information transmission and enhanced price discovery. On WallStreetBets, AI deployment is followed by higher abnormal trading volume, greater volatility, and a higher incidence of lottery-like return outcomes,

consistent with AI amplifying speculative dynamics and sentiment-driven trading. These patterns may not be causal, as unobserved shocks could jointly drive both AI adoption and market outcomes. Nevertheless, the sharply different associations across two platforms with distinct governance and user sophistication provide novel evidence on generative AI and financial information ecosystems.

These findings bear upon policy debates about whether AI deployment enhances or undermines the information ecosystem of financial markets. Our findings suggest that the same tool may facilitate price discovery on a curated platform while amplifying exuberance on an unmoderated one. In essence, the nature of AI-assisted thought transmission—whether it serves to disseminate rational analysis or to engineer persuasive, speculative narratives—is not inherent in the technology itself. Instead, it is critically shaped by the governance, incentives, and community norms of the platform on which it operates.

This highlights a key tension in democratizing financial discourse. While AI can lower barriers to entry, our evidence from WSB suggests that, absent appropriate governance and investor education, there is a risk that AI degrades rather than improves market quality. This risk may be particularly acute in social media environments that act as "product market traps," wherein users participate due to social pressure or a "fear of missing out" (FOMO), even to their own detriment (Bursztyn et al. 2025a,b). Our evidence from financial social media suggests that, in more general social media settings as well, AI presents both dangers and opportunities. Our findings suggest that whether AI proves beneficial or harmful depends on platform protocol design and governance, and on the user expectations about whether the purpose of content is to to entertain, or to express feelings and simplistic beliefs, or to enhance reasoning.

References

Acemoglu, D., A. Ozdaglar, and J. Siderius. 2025. AI and social media: A political economy perspective. Working Paper, Massachusetts Institute of Technology.

Asarch, S. 2021. Reddit's WallStreetBets is facing a culture divide as new users flood the forum. *Business Insider*, https://www.businessinsider.com/wallstreetbets-reddit-forum-divided-as-new-users-flood-subreddit-2021-2.

Bali, T. G., N. Cakici, and R. Whitelaw. 2011. Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics* 99:427–46.

Bali, T. G., D. Hirshleifer, L. Peng, Y. Tang, and Q. Wang. 2025. Social interactions and lottery stock mania. Working Paper, Georgetown University.

Barber, B., and T. Odean. 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance* 55:773–806.

Barber, B. M., X. Huang, T. Odean, and C. Schwarz. 2022. Attention-induced trading and returns: Evidence from Robinhood users. *Journal of Finance* 77:3141–90.

Barrot, J.-N., R. Kaniel, and D. Sraer. 2016. Are retail traders compensated for providing liquidity? *Journal of Financial Economics* 120:146–68.

Bernard, V. L., and J. K. Thomas. 1989. Post-earnings-announcement drift: Delayed price response or risk premium? *Journal of Accounting Research* 27:1–36.

Bertomeu, J., Y. Lin, Y. Liu, and Z. Ni. 2025. The impact of generative AI on information processing: Evidence from the ban of ChatGPT in Italy. *Journal of Accounting and Economics* 80:101782.

Black, F. 1986. Noise. *Journal of Finance* 41:529–43.

Blankespoor, E., E. deHaan, and Q. Li. 2025. Generative AI in financial reporting. Working Paper, University of Washington.

Boehmer, E., C. M. Jones, X. Zhang, and X. Zhang. 2021. Tracking retail investor activity. *Journal of Finance* 76:2249–305.

Bradley, D., J. Hanousek, Jr, R. Jame, and Z. Xiao. 2023. Place your bets? The value of investment research on Reddit's Wallstreetbets. *Review of Financial Studies* 37:1409–59.

Bradshaw, M., H. Ma, A. Yost, and Z. Zou. 2025. Generative AI use by capital-market information intermediaries: Evidence from Seeking Alpha. Working Paper, Boston College.

Brynjolfsson, E., D. Li, and L. Raymond. 2025. Generative AI at work. *Quarterly Journal of Economics* 140:889–942.

Burnside, C., M. Eichenbaum, and S. T. Rebelo. 2016. Understanding booms and busts in housing markets. *Journal of Political Economy* 124:1088–147.

Bursztyn, L., B. R. Handel, R. Jiménez-Durán, and C. Roth. 2025a. When product markets become collective traps: The case of social media. *American Economic Review* 115:4105–36.

Bursztyn, L., A. Imas, R. Jiménez-Durán, A. Leonard, and C. Roth. 2025b. Social dynamics of AI adoption. Working Paper, University of Chicago.

Campbell, J. L., M. D. DeAngelis, and J. R. Moon. 2019. Skin in the game: Personal stock holdings and investors' response to stock analysis on social media. *Review of Accounting Studies* 24:731–79.

Campbell, J. Y., and T. Ramadorai. 2025. Household finance in retrospect and prospect. *Journal of Finance: Insights and Perspectives* Forthcoming.

Cao, S., W. Jiang, J. Wang, and B. Yang. 2024. From man vs. machine to man + machine: The art and AI of stock analyses. *Journal of Financial Economics* 160:103910.

Chan, L. K. C., N. Jegadeesh, and J. Lakonishok. 1996. Momentum strategies. *Journal of Finance* 51:1681–713.

Chang, A., X. Dong, X. Martin, and C. Zhou. 2023. AI democratization and trading inequality. Working Paper, City University of New York, Baruch College.

Chen, H., P. De, Y. J. Hu, and B.-H. Hwang. 2014. Wisdom of crowds: The value of stock opinions transmitted through social media. *Review of Financial Studies* 27:1367–403.

Chordia, T., A. Sarkar, and A. Subrahmanyam. 2005. An empirical analysis of stock and bond market liquidity. *Review of Financial Studies* 18:85–129.

Choukhmane, T., T. de Silva, W. Lin, and M. Akuzawa. 2025. How good is generative ai personal financial advice? Working paper.

Christen, P., D. J. Hand, and N. Kirielle. 2023. A review of the F-Measure: Its history, properties, criticism, and best practice. *ACM Computing Surveys* 56:73.

Condie, E. R., L. L. Lisic, T. A. Seidel, J. M. Truelson, and A. B. Zimmerman. 2023. Does gender and ethnic diversity among audit partners influence office-level audit personnel retention and audit quality? *Contemporary Accounting Research* 40:2477–511. doi: 10.1111/1911-3846.12882.

Cong, L. W., Y. Guo, X. Zhao, and W. Zhou. 2024. Writing quality and soft information in the GenAI

age: Evidence from online credit markets. Working Paper, Cornell University.

Cookson, J. A., J. E. Engelberg, and W. Mullins. 2023. Echo chambers. *Review of Financial Studies* 36:450–500.

Cookson, J. A., R. Lu, W. Mullins, and M. Niessner. 2024. The social signal. *Journal of Financial Economics* 158:103870.

D'Acunto, F., N. R. Prabhala, and A. G. Rossi. 2019. The promises and pitfalls of robo-advising. *The Review of Financial Studies* 32:1983–2020.

Detrixhe, J. 2021. Good luck proving Reddit traders did anything illegal by pumping GameStop. *Quartz*, https://qz.com/1965494/are-wallstreetbets-reddit-traders-manipulating-gamestop-shares.

Diamond, D. W., and R. E. Verrecchia. 1981. Information aggregation in a noisy rational expectations economy. *Journal of Financial Economics* 9:221–35.

Dim, C. 2025. Social media analysts' skill: Evidence from text-implied beliefs. *Journal of Financial and Quantitative Analysis* 60:3081–115.

Dugan, L., A. Hwang, F. Trhlik, J. M. Ludan, A. Zhu, H. Xu, D. Ippolito, and C. Callison-Burch. 2024. RAID: A shared benchmark for robust evaluation of machine-generated text detectors. *arXiv:2405.07940* .

Eaton, G. W., T. C. Green, B. S. Roseman, and Y. Wu. 2022. Retail trader sophistication and stock market quality: Evidence from brokerage outages. *Journal of Financial Economics* 146:502–28.

Farrell, M., T. C. Green, R. Jame, and S. Markov. 2022. The democratization of investment research and the informativeness of retail investor trading. *Journal of Financial Economics* 145:616–41.

Fedyk, A., A. Kakhbod, P. Li, and U. Malmendier. 2025. Ai and perception biases in investments: An experimental study. Working paper.

Financial Stability Board. 2024. The financial stability implications of artificial intelligence.

Glosten, L. R. 1989. Insider trading, liquidity, and the role of the monopolist specialist. *Journal of Business* 62:211–35.

Glosten, L. R., and L. E. Harris. 1988. Estimating the components of the bid/ask spread. *Journal of Financial Economics* 21:123–42.

Glosten, L. R., and P. R. Milgrom. 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14:71–100.

Grossman, S. J., and J. E. Stiglitz. 1980. On the impossibility of informationally efficient markets. *Amer-*

ican Economic Review 70:393–408.

Hackenburg, K., B. M. Tappin, L. Hewitt, E. Saunders, S. Black, H. Lin, C. Fist, H. Margetts, D. G. Rand, and C. Summerfield. 2025. The levers of political persuasion with conversational AI. *arXiv:2507.13919v1* .

Han, B., D. Hirshleifer, and J. Walden. 2022. Social transmission bias and investor behavior. *Journal of Financial and Quantitative Analysis* 57:390–412.

Heimer, R. Z. 2026. Financial decision making in the age of ai. Manuscript in preparation for submission to Oxford Research Encyclopedia of Economics and Finance.

Hirshleifer, D. 2020. Presidential address: Social transmission bias in economics and finance. *Journal of Finance* 75:1779–831.

Hirshleifer, D., S. S. Lim, and S. H. Teoh. 2009. Driven to distraction: Extraneous events and underreaction to earnings news. *Journal of Finance* 64:2289–325.

Hu, D., C. M. Jones, S. Li, V. Zhang, and X. Zhang. 2025. The rise of Reddit: How social media affects belief formation and price discovery. Working Paper, Peking University.

Huang, R. D., and H. R. Stoll. 1997. The components of the bid-ask spread: A general approach. *Review of Financial Studies* 10:995–1034.

Jabarian, B., and A. Imas. 2025. Artificial writing and automated detection. Working Paper, University of Chicago.

Kaniel, R., G. Saar, and S. Titman. 2008. Individual investor sentiment and stock returns. *Journal of Finance* 63:273–310.

Kumar, A., J. K. Page, and O. G. Spalt. 2011. Religious beliefs, gambling attitudes, and financial market outcomes. *Journal of Financial Economics* 102:671–708.

Kyle, A. S. 1985. Continuous auctions and insider trading. *Econometrica* 53:1315–35.

Latona, G. R., M. H. Ribeiro, T. R. Davidson, V. Veselovsky, and R. West. 2024. The AI review lottery: Widespread AI-assisted peer reviews boost paper scores and acceptance rates. *arXiv:2405.02150* .

Lee, D. D., and H. S. Seung. 1999. Learning the parts of objects by non-negative matrix factorization. *Nature* 401:788–91.

Loughran, T., and B. McDonald. 2011. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance* 66:35–65.

———. 2024. Measuring firm complexity. *Journal of Financial and Quantitative Analysis* 59:2487–514.

Lu, J., S. Li, and Y. Shi. 2025. Robo-advising with large language models. *Journal of Finance* Forthcoming.

Lu, Y., N. Y. Naik, and M. Teo. 2024. Diverse hedge funds. *Review of Financial Studies* 37:639–83. doi: 10.1093/rfs/hhad064.

Matz, S. C., J. D. Teeny, S. S. Vaid, H. Peters, G. M. Harari, and M. Cerf. 2024. The potential of generative AI for personalized persuasion at scale. *Scientific Reports* 14:4692.

Merriam-Webster. 2025. Word of the year 2025: Slop.

Mitchell, E., Y. Lee, A. Khazatsky, C. D. Manning, and C. Finn. 2023. DetectGPT: Zero-shot machine-generated text detection using probability curvature. *arXiv:2301.11305* .

Ng, W. J. R., Y. H. M. See, and M. W.-L. Cheung. 2025. The influence of affective and cognitive appeals on persuasion outcomes: A cross-cultural meta-analysis. *Journal of Communication* 75:101–11. doi: 10.1093/joc/jqae042.

Noy, S., and W. Zhang. 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381:187–92.

Pedersen, L. H. 2022. Game on: Social networks and markets. *Journal of Financial Economics* 146:1097–119.

Petty, R. E., and J. T. Cacioppo. 1986. The elaboration likelihood model of persuasion. *Advances in Experimental Social Psychology* 19:123–205.

Pew Research Center. 2025. Social media and news fact sheet.

Rossi, A. G., and S. P. Utkus. 2024. The diversification and welfare effects of robo-advising. *Journal of Financial Economics* 157:103869. doi: 10.1016/j.jfineco.2024.103869.

Schoenegger, P., F. Salvi, J. Liu, X. Nan, R. Debnath, B. Fasolo, E. Leivada, G. Recchia, F. Günther, et al. 2025. Large language models are more persuasive than incentivized human persuaders. *arXiv:2502.09662* .

Shefrin, H., and M. Statman. 1985. The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance* 40:777–90.

Stickel, S. E. 1991. Common stock returns surrounding earnings forecast revisions. *The Accounting Review* 66:402–16.

Stiglitz, J. E., and M. Ventura-Bolet. 2025. The impact of AI and digital platforms on the information ecosystem. Working Paper, Columbia University.

Stoll, H. R. 1989. Inferring the components of the bid-ask spread: Theory and empirical tests. *Journal of Finance* 44:115–34.

Su, J., T. Y. Zhuo, D. Wang, and P. Nakov. 2023. DetectLLM: Leveraging log rank information for zero-shot detection of machine-generated text. *arXiv:2306.05540* .

Wu, J., S. Yang, R. Zhan, Y. Yuan, D. F. Wong, and L. S. Chao. 2025. A survey on LLM-generated text detection: Necessity, methods, and future directions. *Computational Linguistics* 51:275–338.

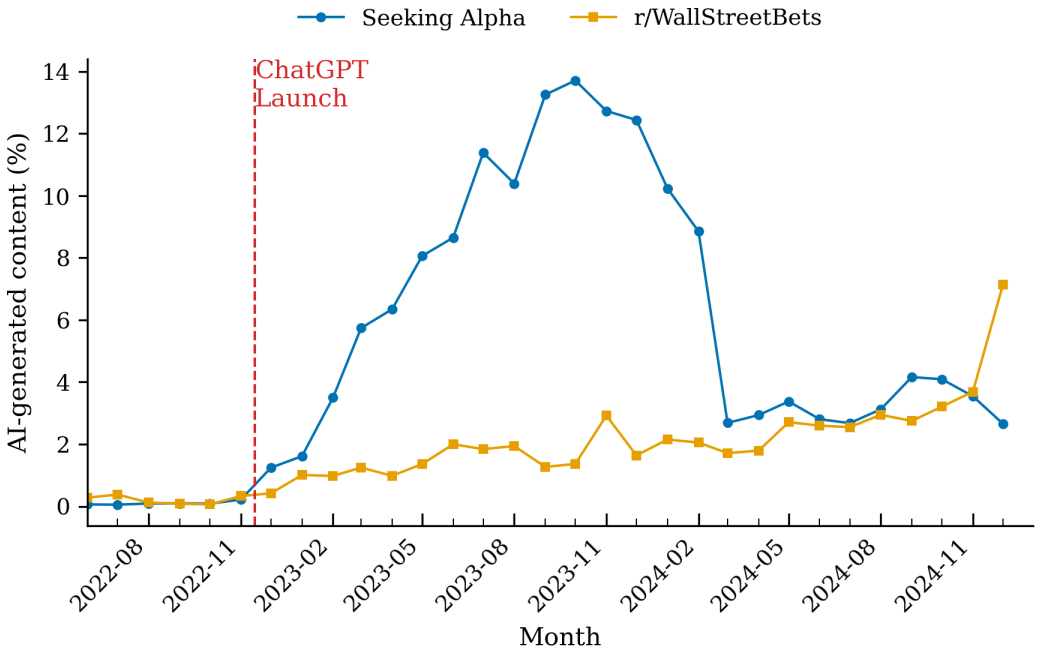


Figure 1: AI-generated Content Over Time

This figure plots the monthly share of AI-generated content in our sample, as classified by GPTZero, on two platforms-Seeking Alpha and Reddit's r/WallStreetBets-from March 2022 through December 2024. The red vertical dashed line marks November 30, 2022, the launch of ChatGPT. The share of AI-generated content on Seeking Alpha declines after the platform begins enforcing its no-AI policy in November 2023. Seeking Alpha further introduced a new submission requirement in May 2024 that further compels authors to verify that their work did not use AI.

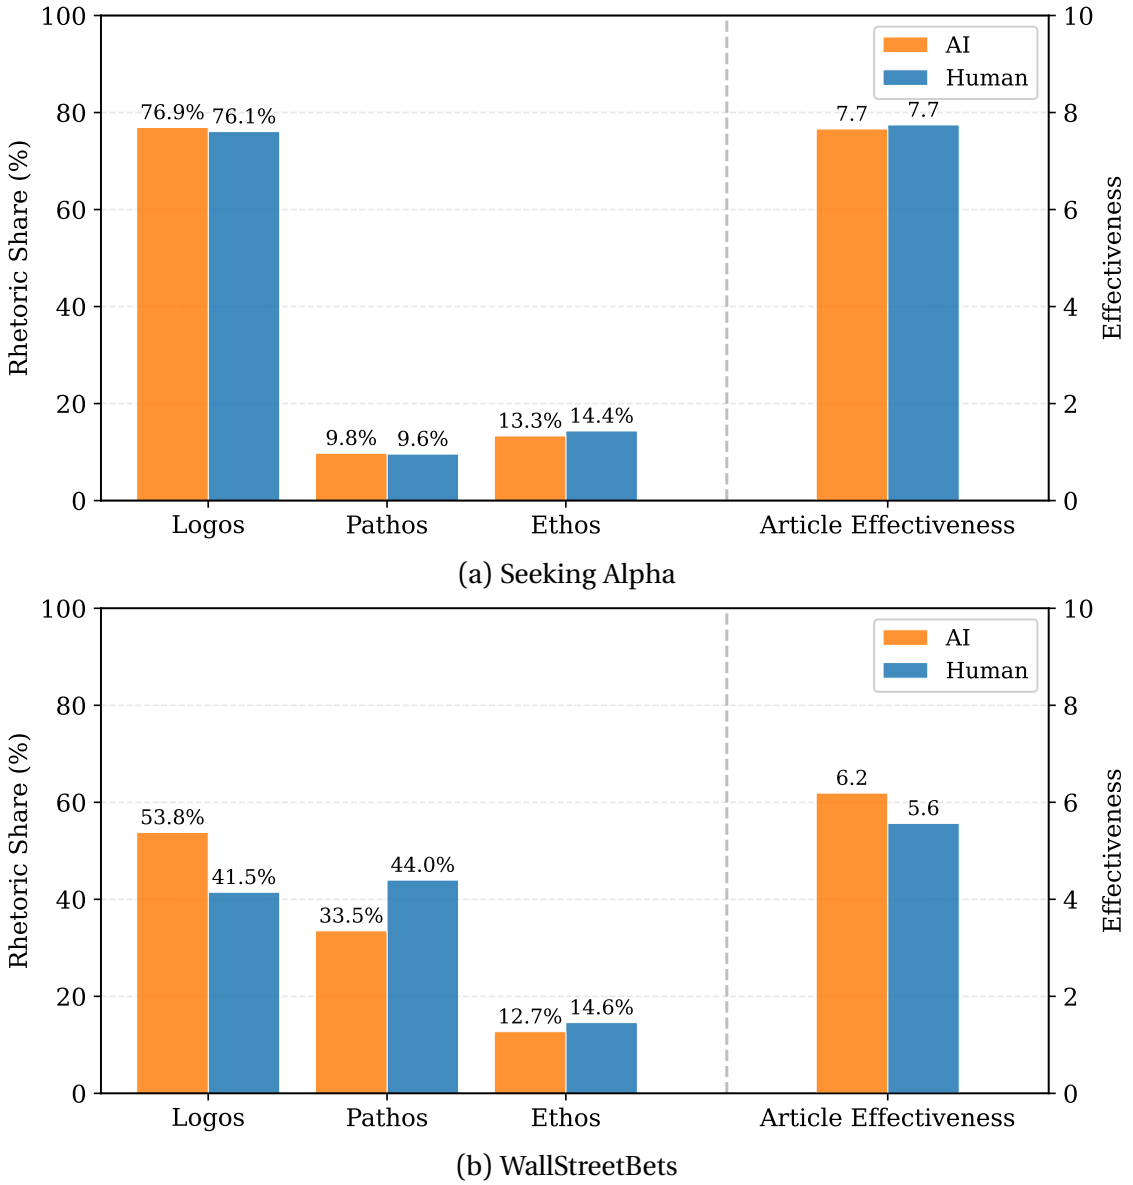


Figure 2: Rhetorical Shares and Article Effectiveness: AI vs Human Content

This figure reports average rhetorical *Share*, defined as the fraction of an article's text devoted to each rhetorical mode: *Logos* (fact- and data-based argumentation), *Pathos* (emotional appeals), and *Ethos* (appeals to credibility and authority), for articles on Seeking Alpha and WallStreetBets, respectively. We use GPT-5.1 to analyze each article and decompose its rhetorical style into three modes, with mode shares summing to one, and to assign an effectiveness score on a 0-10 scale to each mode. Article-level persuasion effectiveness, *Effectiveness*, is the sum across modes of each mode's share multiplied by its effectiveness. The left y-axis reports rhetorical shares, and the right y-axis reports effectiveness scores. Orange bars represent AI-generated content; red bars represent human-authored content.

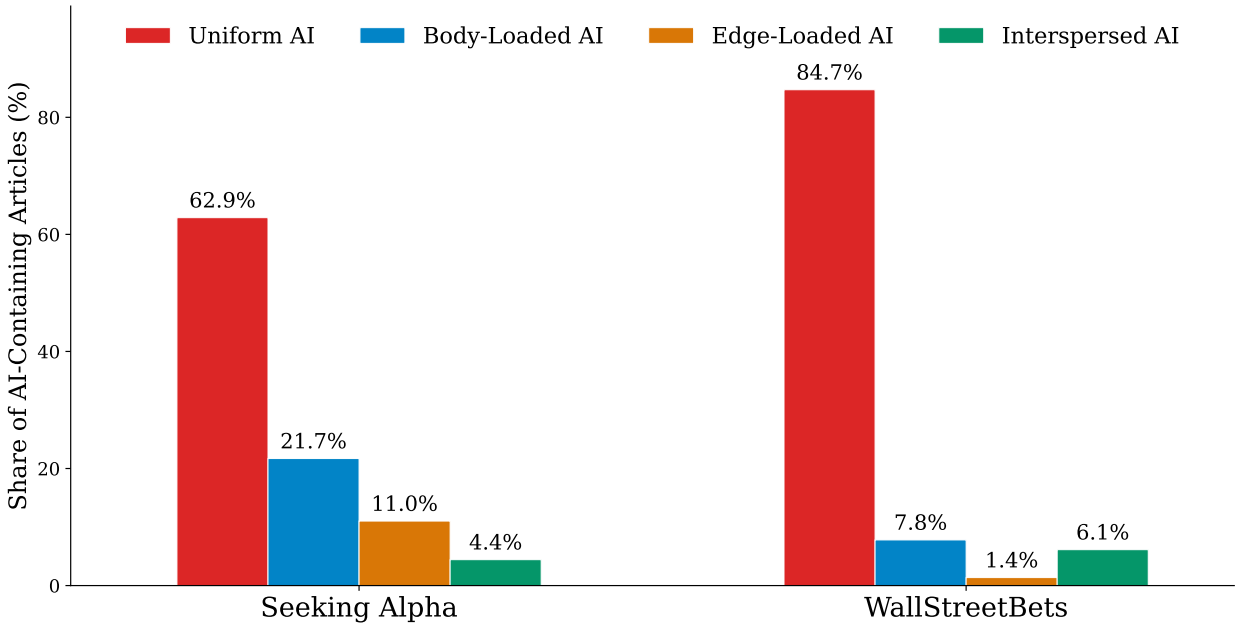


Figure 3: AI Usage Category Distribution by Platform.

This figure shows the share of articles classified into each usage category on Seeking Alpha (SA) and WallStreetBets (WSB). The sample is restricted to articles classified as "AI-ONLY" or "MIXED" by GPTZero, excluding articles with fewer than five sentences. Categories are defined based on the positional distribution of sentence-level AI probability within each article. See Section 7.1 for category definitions and Internet Appendix for classification methodology.

Table 1. Summary statistics

This table presents summary statistics of article, author and firm characteristics from Seeking Alpha and Reddit's r/WallStreetBets during the sample period (December 2022–December 2024). Panel A reports summary statistics on average article characteristics; Panel B, on total author activity; and Panel C, on variables at the stock-day level. Detailed definitions of all variables are provided in Appendix Table A1.

Panel A: Article Characteristics

	SA						WSB					
	N	Mean	SD	p25	p50	p75	N	Mean	SD	p25	p50	p75
AI Prob (%)	57,581	10.65	21.65	0.91	2.75	6.28	64,848	3.60	14.36	0.52	0.68	1.57
Length	57,581	1453.67	675.77	1022	1309	1705	64,848	114.18	124.75	61	79	119
Sentiment	57,581	0.03	0.01	0.02	0.02	0.03	64,848	0.02	0.03	0.00	0.02	0.04
Complexity (%)	57,581	0.22	0.27	0.05	0.14	0.31	64,848	0.09	0.37	0.00	0.00	0.00
Fog Index	57,581	14.57	2.38	12.96	14.39	15.88	64,848	10.66	6.93	7.59	9.48	12.00
Quantitative (%)	57,581	3.80	2.24	2.38	3.41	4.73	64,848	2.67	3.52	0.00	1.69	3.85
Graphical (%)	57,581	0.46	0.31	0.24	0.41	0.62	64,848	0.02	0.25	0.00	0.00	0.00
Comments	57,581	11.32	20.00	1	4	13	64,848	3.85	36.67	0	1	2
Disagreement	40,547	0.06	0.04	0.03	0.05	0.08	21,368	0.06	0.06	0.02	0.04	0.08

Panel B: Author Characteristics

	SA						WSB					
	N	Mean	SD	p25	p50	p75	N	Mean	SD	p25	p50	p75
Author Article Count	1,998	28.78	80.88	2	5	18	28,563	2.27	5.07	1	1	2
# Covered Stocks	1,998	16.24	36.90	1	4	13	28,563	1.64	2.09	1	1	2
# Covered Industries	1,998	10.24	18.92	1	4	10	28,563	1.60	1.77	1	1	2

Panel C: Stock-Day Sample Characteristics

	SA						WSB					
	N	Mean	SD	p25	p50	p75	N	Mean	SD	p25	p50	p75
AIDay (×100)	1,607,466	0.24	4.85	0	0	0	687,418	0.18	4.28	0	0	0
Articles (×100)	1,607,466	3.59	22.37	0	0	0	687,418	9.43	136.15	0	0	0
Sentiment (×100)	1,607,466	0.08	0.50	0	0	0	687,418	0.06	0.60	0	0	0
AR_t (%)	1,459,930	0.02	5.18	-1.25	-0.04	1.15	648,718	0.03	4.31	-1.17	-0.04	1.08
$CAR_{[t+1,t+5]}$ (%)	1,594,720	0.11	11.68	-3.46	-0.18	3.07	682,020	0.13	12.65	-3.31	-0.18	2.92
$AVOL_{[t+1,t+5]}$ (%)	1,591,465	2.87	60.09	-29.55	-2.29	28.54	680,946	3.02	56.58	-27.32	-2.10	26.51
$Volatility_{[t+1,t+5]}$ (%)	1,594,720	2.85	4.64	1.31	2.09	3.42	682,020	2.78	3.94	1.24	1.98	3.28
$Spread_{[t+1,t+5]}$ (%)	1,520,352	0.41	1.21	0.07	0.14	0.41	653,059	0.29	0.60	0.05	0.09	0.21
Size ($B)	1,484,486	14.11	91.72	0.30	1.35	5.32	653,497	26.67	135.58	0.76	3.01	13.88
BM	1,484,486	0.62	0.96	0.22	0.45	0.80	653,497	0.56	0.65	0.19	0.40	0.73
IO	1,484,486	0.68	0.35	0.50	0.77	0.90	653,497	0.71	0.35	0.60	0.80	0.90
Analysts	1,484,486	1.91	1.86	0.37	1.10	2.58	653,497	2.46	2.19	0.74	1.84	3.68
News	1,484,486	2.59	1.28	1.79	2.71	3.43	653,497	2.92	1.32	2.20	3.09	3.76
Earnings Day	1,484,486	0.01	0.12	0	0	0	653,497	0.02	0.12	0	0	0
$Ret_{[t-21,t-1]}$ (%)	1,484,486	1.60	23.80	-6.94	0.30	7.95	653,497	1.72	20.35	-6.56	0.47	7.75
$Volatility_{[t-21,t-1]}$ (%)	1,605,333	3.19	4.32	1.68	2.47	3.83	686,727	3.12	3.52	1.59	2.33	3.69

Table 2. Author Characteristics and AI Adoption on SeekingAlpha

This table examines how an article's AI probability score relates to author characteristics on Seeking Alpha using the following regression:

$$AIProb_{ijt} = \beta_1 AuthorChars_j + Controls + FE + \epsilon_{ijt},$$

where the dependent variable $AIProb_{ijt}$ is the AI probability score (in percentage points) of article j by covering stock i on day t. $AuthorChars_j$ are indicator variables that correspond to an given author characteristic. We control for article characteristics, including length, sentiment, complexity, fog index, graphical, quantitative, and a dummy variable for submissions on WSB. We also include day and stock fixed effects. Columns 1–8 report results from specifications that include one of the author characteristics; Column 9 reports the full specification that includes all characteristics jointly. Standard errors are two-way clustered by stock and day; t-statistics are in parentheses. *p <.1; **p <.05; ***p <.01.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
English as Primary Lang.	-2.002***								-1.469***
	(-5.76)								(-4.14)
Certificated		-1.775***							-1.773***
		(-4.64)							(-4.39)
Portfolio Manager			-3.692***						-4.215***
			(-9.26)						(-10.12)
Trader				-0.714**					-0.206
				(-2.08)					(-0.60)
Advisor					-1.973***				-1.714***
					(-6.07)				(-5.37)
Analyst						4.475***			4.977***
						(8.03)			(8.67)
PhD							7.767***		8.596***
							(9.05)		(9.55)
MBA								1.347**	1.894***
								(2.51)	(3.42)
No Bio	0.005	0.138	0.044	0.155	0.028	0.582	0.528	0.339	0.146
	(0.01)	(0.25)	(0.08)	(0.27)	(0.05)	(1.02)	(0.94)	(0.60)	(0.25)
Article Ctrl.	X	X	X	X	X	X	X	X	X
Day FE	X	X	X	X	X	X	X	X	X
Stock FE	X	X	X	X	X	X	X	X	X
Obs.	57,581	57,581	57,581	57,581	57,581	57,581	57,581	57,581	57,581
Adj. R^2 (%)	10.0	9.9	10.0	9.9	10.0	10.3	10.2	9.9	11.0

Table 3. Author Unfamiliarity, Information Scarcity and AI Adoption

This table examines whether AI use is associated with author unfamiliarity (Panel A) and information scarcity, proxied by recent news flow (Panel B), at the article level. We estimate the same equation as in Table 2, replacing the key explanatory variable with measures of unfamiliarity or information scarcity. $AIProb_{ijt}$ (in percentage points) is the AI probability score for content j about firm i, posted on day t. $Unfamiliarity_{ijt}$ equals to one if the author has not covered firm i in the prior six months. $News_{it}$ is the log of one plus the number firm-specific news items in the prior three days and $Analysts_{it}$ is the log of one plus the number of analysts covering the firm in the most recent quarter. Both $Newsit$ and $Analystsit$ are standardized to have zero mean and unit variance. *Controls* include firm characteristics and measures of the author's prior activity and performance (each standardized to zero mean and unit variance). Columns 1–2 in each panel report results for articles published on Seeking Alpha; Columns 3–4 in each panel report results for posts on Reddit's r/WallStreetBets. Standard errors are two-way clustered by stock and day; t-statistics are shown in parentheses. *p <.1; **p <.05; ***p <.01.

Panel A: Author Unfamiliarity

	SA		WSB	
	(1)	(2)	(3)	(4)
Unfamiliarity	2.322***	0.555***	0.007	-0.012
	(9.42)	(3.26)	(0.04)	(-0.08)
Size	-0.573	-1.781**	-1.398*	-0.527
	(-0.54)	(-2.20)	(-1.86)	(-0.54)
BM	0.919***	0.291	0.162	-0.320
	(2.84)	(0.86)	(1.00)	(-1.40)
IO	0.943*	0.452	0.363	-0.480
	(1.69)	(1.21)	(0.95)	(-1.19)
Author Article	-3.418***	-1.109***	-0.808***	-0.055
	(-20.60)	(-8.10)	(-8.50)	(-0.58)
Author AI Article	8.358***	2.497***	2.356***	-0.431***
	(36.41)	(19.15)	(6.99)	(-2.68)
Author Comments	0.490***	0.468***	-0.168*	0.061
	(3.32)	(3.74)	(-1.78)	(0.56)
Submissions			9.089***	4.364***
			(10.95)	(5.88)
Day FE	X	X	X	X
Stock FE	X	X	X	X
Author FE		X		X
Obs.	54,429	54,429	63,400	63,400
Adj. R^2 (%)	18.7	59.5	11.5	56.6

Panel B: Information Scarcity				
	SA		WSB	
	(1)	(2)	(3)	(4)
News	-0.319**	-0.206**	-0.050	-0.215
	(-2.15)	(-2.11)	(-0.37)	(-1.41)
Analysts	-0.729***	-0.332**	0.527**	0.162
	(-3.42)	(-1.98)	(2.42)	(0.79)
Size	-0.389	-1.673**	-1.727**	-0.520
	(-0.37)	(-2.06)	(-2.06)	(-0.50)
BM	0.920***	0.288	0.221	-0.326
	(2.81)	(0.84)	(1.34)	(-1.39)
IO	0.960*	0.454	0.429	-0.463
	(1.73)	(1.22)	(1.14)	(-1.18)
Author Article	-3.678***	-1.131***	-0.809***	-0.053
	(-22.16)	(-8.26)	(-9.29)	(-0.59)
Author AI Article	8.383***	2.497***	2.356***	-0.430***
	(36.57)	(19.16)	(6.99)	(-2.67)
Author Comments	0.479***	0.472***	-0.169*	0.061
	(3.22)	(3.76)	(-1.78)	(0.56)
Submissions			9.088***	4.363***
			(10.93)	(5.88)
Day FE	X	X	X	X
Stock FE	X	X	X	X
Author FE		X		X
Obs.	54,429	54,429	63,400	63,400
Adj. R^2 (%)	18.5	59.5	11.5	56.6

Table 4. Intense Retail Buying and AI Adoption

This table examines the association between recent intense retail buying and article-level AI usage using the following panel regression specification:

$$AIProb_{ijt} = \beta_1 IntenseRetailBuy_{i[t-3,t-1]} + Controls + FE + \epsilon_{ijt},$$

where $AIProb_{ijt}$ (in percentage points) is the AI probability score for content j about firm i, posted on day t. *IntenseRetailBuy* is an indicator equal to one if retail order imbalance over [t-3, t-1] ranks in cross-sectional top 10%. *Controls* include firm characteristics and measures of the author's prior activity/performance (each standardized to zero mean and unit variance). Columns 1–2 report results for articles published on Seeking Alpha; Columns 3–4 report results for posts on Reddit's r/WallStreetBets. Standard errors are two-way clustered by stock and day; t-statistics are in parentheses. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

	SA		WSB	
	(1)	(2)	(3)	(4)
IntenseRetailBuy	-0.186	-0.460	1.588**	1.992***
	(-0.31)	(-1.10)	(2.53)	(3.19)
Size	-0.775	-1.844**	-1.548**	-0.713
	(-0.72)	(-2.23)	(-2.00)	(-0.68)
BM	0.949***	0.347	0.227	-0.226
	(2.85)	(1.00)	(1.47)	(-1.12)
IO	1.095**	0.440	0.427	-0.565
	(1.99)	(1.13)	(1.07)	(-1.37)
Author Article	-3.719***	-1.135***	-0.606***	-0.043
	(-21.86)	(-8.04)	(-9.01)	(-0.49)
Author AI Article	8.499***	2.510***	0.757***	-0.362**
	(36.76)	(18.87)	(5.97)	(-2.52)
Author Comments	0.428***	0.493***	0.031	0.071
	(2.80)	(3.83)	(0.39)	(0.61)
Submissions			8.812***	4.320***
			(10.62)	(5.50)
Day FE	X	X	X	X
Stock FE	X	X	X	X
Author FE		X		X
Obs.	52,695	52,695	60,473	60,473
Adj. R^2 (%)	18.8	60.0	9.6	57.9

Table 5. AI-Generated Content and Sentiment Alignmment

This table examines the association between an article's AI probability score and the alignment between article sentiment and subsequent comment sentiment using the following regression:

$$CommentSentiment_{ij[t,t+10]} = \beta_1 AIProb_{ijt} \times Sentiment_{ijt} + \beta_2 Sentiment_{ijt} + \beta_3 AIProb_{ijt} + Controls + FE + \epsilon_{ijt}$$

where $CommentSentiment_{ij[t,t+10]}$ is the mean sentiment of comments following content j about firm i over $[t, t+10]$. $AIProb_{ijt}$ is the estimated probability, ranging from 0 to 1, that content j covering firm i on day t is AI-generated. $Sentiment_{ijt}$ is the main text sentiment score. *Controls* includes the article characteristics, firm characteristics, and author prior activities (standardized to zero mean and unit variance). Columns 1–2 report results for articles published on Seeking Alpha; Columns 3–4 report results for posts on Reddit's r/WallStreetBets. Standard errors are two-way clustered by stock and day; t-statistics are in parentheses. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

	SA		WSB	
	(1)	(2)	(3)	(4)
AIProb × Sentiment	-0.022	-0.046	0.090**	0.135**
	(-0.40)	(-0.83)	(2.38)	(2.22)
Sentiment	0.227***	0.217***	0.080***	0.067***
	(13.49)	(11.85)	(12.93)	(7.27)
AIProb	-0.001	-0.001	-0.004**	-0.008**
	(-0.80)	(-0.50)	(-2.49)	(-2.43)
Length	0.003***	0.002***	0.000	0.000
	(14.10)	(7.17)	(-1.61)	(0.66)
Complexity	0.000*	0.000	0.000	0.000
	(-1.95)	(-0.06)	(-0.06)	(-0.17)
Fog Index	0.000	0.001***	0.000	0.000
	(-1.41)	(3.46)	(-1.58)	(0.29)
Graphical	0.002***	0.001***	0.000	0.000
	(8.79)	(4.32)	(0.63)	(0.03)
Quantitative	0.001***	0.000	0.001***	0.001**
	(2.81)	(0.69)	(3.98)	(2.20)
Size	0.002	0.003	-0.001	0.000
	(0.85)	(1.25)	(-0.16)	(-0.03)
BM	-0.001	0.000	-0.001	0.000
	(-0.93)	(-0.66)	(-1.59)	(-0.11)
IO	0.001	0.001	0.000	0.001
	(0.94)	(0.57)	(0.06)	(0.37)
News	0.000	0.000	0.000	0.000
	(0.38)	(-0.11)	(-0.62)	(-0.49)
Analysts	0.000	0.000	0.000	0.002
	(0.20)	(0.00)	(0.93)	(1.57)
Author Article	-0.001***	0.000	0.000	0.000
	(-3.46)	(-1.32)	(0.24)	(-0.24)
Author AI Article	0.000	0.000	0.000	0.000
	(1.14)	(0.01)	(-0.20)	(0.18)
Author Comments	0.001***	0.000	0.000	0.000
	(5.11)	(1.01)	(1.32)	(-0.61)
Submissions			0.021***	0.019***
			(33.19)	(11.23)
Day FE	X	X	X	X
Stock FE	X	X	X	X
Author FE		X		X
Obs.	54,429	54,429	63,400	63,400
Adj. R^2 (%)	8.3	10.1	3.4	0.1

Table 6. AI Content and Return Predictability

This table examines whether AI content can better predict future returns using the following article-level regression:

$$\text{CAR}_{i[t+1,t+5]} = \beta_1 AIProb_{ijt} \times Sentiment_{ijt} + \beta_2 Sentiment_{ijt} + \beta_3 AIProb_{ijt} + Controls + FE + \epsilon_{ijt}.$$

The dependent variable is cumulative abnormal returns over $[t+1, t+5]$ in columns 1 and 3, and over $[t+6, t+10]$ in columns 2 and 4. $AIProb_{ijt}$ is the estimated probability, ranging from 0 to 1, that content j covering firm i on day t is AI-generated. Sentiment is the article-level sentiment score. All continuous control variables are standardized. Standard errors are two-way clustered by date and firm. t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	SA		WSB	
	$\text{CAR}_{[t+1,t+5]}$ (1)	$\text{CAR}_{[t+6,t+10]}$ (2)	$\text{CAR}_{[t+1,t+5]}$ (3)	$\text{CAR}_{[t+6,t+10]}$ (4)
AIProb × Sentiment	0.350**	-0.122	-0.179	0.007
	(2.114)	(-0.851)	(-0.793)	(0.034)
Sentiment	-0.149***	-0.075*	0.019	-0.020
	(-3.398)	(-1.891)	(0.500)	(-0.451)
AIProb	-0.327*	-0.093	-0.259	0.077
	(-1.675)	(-0.579)	(-1.009)	(0.262)
Length	0.070*	-0.005	0.006	0.022
	(1.806)	(-0.145)	(0.167)	(0.681)
Complexity	-0.019	-0.002	-0.019	-0.013
	(-0.663)	(-0.075)	(-0.490)	(-0.365)
Fog Index	0.014	0.021	-0.100**	-0.039
	(0.399)	(0.645)	(-2.355)	(-1.099)
Graphical	0.087**	0.059*	-0.061*	-0.046
	(2.440)	(1.952)	(-1.903)	(-1.214)
Quantitative	0.035	0.026	-0.054	-0.056
	(1.112)	(1.024)	(-1.607)	(-1.343)
Size	-3.549***	-2.927***	-2.857*	-3.197**
	(-3.959)	(-3.716)	(-1.801)	(-2.020)
BM	-0.484	-0.329	0.327	-0.138
	(-1.610)	(-1.300)	(0.707)	(-0.380)
Analysts	-0.124	-0.041	0.319	0.310
	(-1.144)	(-0.366)	(0.626)	(0.990)
IO	-1.002**	-0.624**	-0.120	0.376
	(-2.195)	(-2.443)	(-0.260)	(0.914)
News	-0.089	-0.034	-0.299	-0.091
	(-0.759)	(-0.343)	(-0.680)	(-0.191)
Earnings Day	0.156	-0.119	-0.619	-0.211
	(0.675)	(-0.436)	(-1.073)	(-0.344)
$\text{Ret}_{[t-21,t-1]}$	-0.189	-0.309***	-0.806**	0.305
	(-0.905)	(-2.792)	(-2.155)	(0.629)
$\text{Volatility}_{[t-21,t-1]}$	-0.010	0.202	0.247	0.057
	(-0.058)	(1.620)	(0.595)	(0.135)
Submissions			0.205	0.250
			(1.580)	(1.280)
Day FE	Yes	Yes	Yes	Yes
Stock FE	Yes	Yes	Yes	Yes
Obs.	54,070	54,070	62,785	62,785
Adj. R^2 (%)	8.6	9.1	20.0	19.3

Table 7. AI Day and Retail Order Flow Informativeness

This table examines the association between "AI Day" and the relationship between retail order flow and subsequent returns, using the following stock-day panel regression specification:

$$\begin{aligned}\text{CAR}_{i[t+1,t+5]} &= \beta_1 AIDay_{it} \times \text{OIB}_{it} + \beta_2 HumanDay_{it} \times \text{OIB}_{it} \\ &+ \beta_3 AIDay_{it} \times Sentiment_{it} + \beta_4 HumanDay_{it} \times Sentiment_{it} \\ &+ \beta_5 AIDay_{it} + \beta_6 HumanDay_{it} + \beta_7 \text{OIB} + \beta_8 Sentiment_{it} \\ &+ \beta_9 Attention_{it} + Controls + FE + \epsilon_{it},\end{aligned}$$

where $\text{CAR}_{[t+1,t+5]}$ is the market-adjusted cumulative abnormal return over $[t+1, t+5]$. $\text{OIB}_{i,t}$ is the net retail order imbalance (net retail buys divided by total retail share volume) for stock i on day t. $AIDay_{it}^{SA}$ is an indicator equal to one if at least one AI-generated article/post about stock i appears on day t on SA. $AIDay_{it}^{WSB}$ is defined similarly for WSB. $HumanDay_{it}^{SA}$ is an indicator equal to one if at least one Human article/post about stock i appears on day t on SA. $Sentiment_{it}^{SA}$ is the average platform-level sentiment across all content about stock i on day t on SA. $Attention_{it}^{SA}$ is the logarithm of one plus the number of articles about stock i on day t on SA. All regressions include firm-characteristic controls and stock and day fixed effects. Standard errors are two-way clustered by stock and day. t-statistics are reported in parentheses. Significance levels are denoted as: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	SA		WSB		Combined	
	(1)	(2)	(3)	(4)	(5)	(6)
$AIDay^{SA} \times OIB$	0.403*	0.431**			0.407*	0.435**
	(1.940)	(2.069)			(1.929)	(2.057)
$HumanDay^{SA} \times OIB$	0.053	0.053			0.050	0.050
	(0.795)	(0.793)			(0.749)	(0.747)
$AIDay^{SA}$	-0.294	-0.577			-0.297	-0.573
	(-1.549)	(-1.565)			(-1.596)	(-1.564)
$HumanDay^{SA}$	0.032	0.086			0.027	0.081
	(0.180)	(0.456)			(0.156)	(0.433)
$Sentiment^{SA}$	-0.021	0.012			-0.021	0.012
	(-1.593)	(0.305)			(-1.559)	(0.306)
$Attention^{SA}$	0.019	0.015			0.019	0.015
	(0.595)	(0.456)			(0.621)	(0.476)
$AIDay^{SA} \times Sentiment^{SA}$		0.025				0.023
		(0.398)				(0.380)
$HumanDay^{SA} \times Sentiment^{SA}$		-0.041				-0.040
		(-0.927)				(-0.914)
$AIDay^{WSB} \times OIB$			0.371	0.325	0.478	0.478
			(0.998)	(0.884)	(1.061)	(1.060)
$HumanDay^{WSB} \times OIB$			0.070	0.069	0.090	0.090
			(0.954)	(0.945)	(0.995)	(0.991)
$AIDay^{WSB}$			0.064	-0.097	-0.086	-0.085
			(0.195)	(-0.269)	(-0.235)	(-0.234)
$HumanDay^{WSB}$			0.183	0.194	0.189	0.188
			(1.235)	(1.317)	(1.251)	(1.241)
$Sentiment^{WSB}$			-0.002	-0.139	-0.090	-0.091
			(-0.224)	(-1.094)	(-1.070)	(-1.072)
$Attention^{WSB}$			-0.028	-0.029	-0.019	-0.019
			(-0.773)	(-0.797)	(-0.753)	(-0.741)
$AIDay^{WSB} \times Sentiment^{WSB}$				0.121	0.077	0.077
				(1.004)	(0.953)	(0.954)
$HumanDay^{WSB} \times Sentiment^{WSB}$				0.136	0.090	0.090
				(1.082)	(1.065)	(1.067)
OIB	0.004	0.004	0.022	0.022	-0.003	-0.003
	(0.174)	(0.174)	(1.021)	(1.024)	(-0.132)	(-0.132)
Size	-2.943***	-2.943***	-2.945***	-2.945***	-3.002***	-3.002***
	(-10.493)	(-10.493)	(-9.177)	(-9.179)	(-10.600)	(-10.600)
BM	-0.295***	-0.295***	-0.307**	-0.307**	-0.281***	-0.281***
	(-2.721)	(-2.721)	(-2.312)	(-2.313)	(-2.621)	(-2.620)
IO	-0.046	-0.046	-0.056	-0.056	-0.007	-0.007
	(-1.060)	(-1.060)	(-0.632)	(-0.632)	(-0.101)	(-0.101)
Analysts	-0.166***	-0.166***	-0.095*	-0.095*	-0.172***	-0.172***
	(-3.026)	(-3.026)	(-1.831)	(-1.830)	(-3.128)	(-3.128)
News	-0.075	-0.075	-0.145**	-0.145**	-0.087*	-0.087*
	(-1.616)	(-1.617)	(-2.381)	(-2.381)	(-1.893)	(-1.893)
Earnings Day	-0.103	-0.102	-0.080	-0.081	-0.120	-0.119
	(-0.918)	(-0.913)	(-0.664)	(-0.668)	(-1.064)	(-1.059)
$Ret_{[t-21,t-1]}$	-0.397***	-0.397***	-0.342***	-0.342***	-0.433***	-0.433***
	(-9.301)	(-9.300)	(-5.117)	(-5.118)	(-9.613)	(-9.613)
$Volatility_{[t-21,t-1]}$	0.060	0.060	0.076	0.076	0.056	0.056
	(0.695)	(0.695)	(0.727)	(0.727)	(0.716)	(0.716)
Day FE	X	X	X	X	X	X
Stock FE	X	X	X	X	X	X
Obs.	1,370,181	1,370,181	611,883	611,883	1,391,849	1,391,849
Adj. R^2 (%)	3.2	3.2	4.4	4.4	3.2	3.2

Table 8. AI Day and Subsequent Volume, Volatility, and Bid-Ask Spreads

This table examines the association between platform-level "AI Day" and subsequent trading metrics over $[t+1, t+5]$, using the following stock-day panel regression specification:

$$MarketQuality_{i[t+1,t+5]} = \beta_1 AIDay_{it} + \beta_2 HumanDay_{it} + \beta_3 Attention_{it} + \beta_4 Sentiment_{it} + Controls + FE + \epsilon_{it},$$

$MarketQuality_{i[t+1,t+5]}$ measures trading activity for stock i, standardized to zero mean and unit variance. Columns (1), (4), and (7) report abnormal daily volume; columns (2), (5), and (8) report return volatility; and columns (3), (6), and (9) report the effective bid-ask spread. $AIDay_{it}^{SA}$ is an indicator equal to one if at least one AI-generated article/post about stock i appears on day t on SA. $HumanDay_{it}^{SA}$ is an indicator equal to one if at least one Human article/post about stock i appears on day t on SA. $Sentiment_{it}^{SA}$ is the average platform-level sentiment across all content about stock i on day t on SA. $Attention_{it}^{SA}$ is the logarithm of one plus the number of articles about stock i on day t on SA. $AIDay_{it}^{WSB}$, $HumanDay_{it}^{WSB}$, $Sentiment_{it}^{WSB}$, and $Attention_{it}^{WSB}$ are defined similarly for WSB. All regressions include firm-characteristic controls and stock and day fixed effects. Standard errors are two-way clustered by stock and day. t-statistics are reported in parentheses. Significance levels are denoted as: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	SA			WSB			Combined		
	AVOL	Volatility	Spread	AVOL	Volatility	Spread	AVOL	Volatility	Spread
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$AIDay^{SA}$	-0.023	-0.019	-0.018***				0.015	0.022	-0.016***
	(-1.02)	(-0.92)	(-5.71)				(0.72)	(1.16)	(-5.45)
$HumanDay^{SA}$	0.006	0.017	-0.009***				0.046**	0.060***	-0.007***
	(0.27)	(0.99)	(-3.08)				(2.19)	(3.45)	(-2.65)
$Attention^{SA}$	0.020***	0.013***	0.001**				0.011***	0.003	0.001
	(4.90)	(4.00)	(2.15)				(2.81)	(1.08)	(1.51)
$Sentiment^{SA}$	-0.004***	-0.012***	0.000				-0.003**	-0.011***	0.000
	(-2.81)	(-8.74)	(-0.30)				(-2.37)	(-8.61)	(-0.17)
$AIDay^{WSB}$				0.086***	0.074*	0.000	0.088***	0.078*	0.004
				(2.82)	(1.87)	(-0.03)	(2.96)	(1.96)	(0.85)
$HumanDay^{WSB}$				-0.017	-0.066***	-0.004	-0.014	-0.062***	-0.004
				(-0.82)	(-2.96)	(-1.30)	(-0.67)	(-2.75)	(-1.41)
$Attention^{WSB}$				0.038***	0.045***	0.002***	0.024***	0.028***	0.001***
				(6.50)	(6.42)	(2.93)	(5.69)	(6.15)	(2.60)
$Sentiment^{WSB}$				0.006***	0.003**	0.000	0.004***	0.002***	0.000
				(5.32)	(2.55)	(-1.55)	(5.34)	(2.60)	(-0.98)
Size	-0.192***	-0.178***	-0.111***	-0.165***	-0.193***	-0.100***	-0.198***	-0.180***	-0.105***
	(-6.37)	(-6.80)	(-8.67)	(-4.57)	(-5.45)	(-4.69)	(-6.78)	(-7.02)	(-8.48)
BM	-0.007	-0.042***	0.005	-0.025*	-0.022*	-0.001	-0.009	-0.037***	0.006
	(-0.71)	(-3.99)	(1.15)	(-1.84)	(-1.66)	(-0.12)	(-0.93)	(-3.74)	(1.24)
IO	-0.018***	-0.006	-0.005*	-0.008***	-0.003	-0.007	-0.016**	-0.004	-0.007*
	(-2.63)	(-0.66)	(-1.72)	(-3.03)	(-1.10)	(-1.17)	(-2.56)	(-0.49)	(-1.82)
Analysts	-0.036***	-0.003	-0.002	-0.036***	-0.003	-0.003	-0.035***	-0.003	-0.002
	(-7.16)	(-0.61)	(-0.97)	(-4.99)	(-0.48)	(-1.32)	(-7.13)	(-0.51)	(-1.09)
News	0.063***	-0.057***	-0.024***	0.032***	-0.080***	-0.014***	0.063***	-0.056***	-0.024***
	(8.23)	(-8.32)	(-12.97)	(3.15)	(-11.03)	(-5.23)	(7.95)	(-8.23)	(-12.77)
Earnings Day	0.430***	0.140***	0.003	0.465***	0.114***	0.005**	0.413***	0.134***	0.002
	(29.52)	(9.24)	(1.25)	(28.63)	(7.82)	(2.14)	(29.62)	(9.32)	(0.75)
$Ret_{[t-21,t-1]}$	0.255***	0.020***	-0.022***	0.316***	0.018**	-0.029***	0.258***	0.021***	-0.022***
	(4.24)	(2.82)	(-3.12)	(25.74)	(2.14)	(-10.31)	(4.40)	(2.92)	(-3.22)
$Volatility_{[t-21,t-1]}$	0.028	0.047	-0.011***	0.116***	0.099***	-0.021***	0.036	0.051	-0.012**
	(0.76)	(1.51)	(-2.85)	(3.83)	(5.08)	(-3.66)	(0.89)	(1.59)	(-2.57)
$Spread_{[t-21,t-1]}$			0.644***			0.605***			0.643***
			(41.36)			(15.78)			(41.71)
Day FE	X	X	X	X	X	X	X	X	X
Stock FE	X	X	X	X	X	X	X	X	X
Obs.	1,472,523	1,472,523	1,343,895	648,295	648,295	594,458	1,496,823	1,496,823	1,365,371
Adj. R^2 (%)	17.6	39.0	88.6	20.8	42.1	90.6	17.3	38.9	88.6

Table 9. Lottery-Like Return Events Following AI Days

This table examines the association between "AI Day" and the incidence of large-return events (MAX and Lottery) using the following logistic stock-day panel regression specification:

$$\begin{aligned}\text{MAX/}Lottery_{i[t+1,t+5]} &= \beta_1 AIDay_{it} + \beta_2 HumanDay_{it} + \beta_3 Attention_{it} + \beta_4 Sentiment_{it} \\ &\quad + Controls + FE + \epsilon_{it},\end{aligned}$$

where the dependent variables, $MAX_{i[t+1,t+5]}$ and $Lottery_{i[t+1,t+5]}$, are indicators equal to one if stock i experiences a MAX or Lottery event in the $[t+1, t+5]$ window, and zero otherwise. A MAX event occurs when the daily return is the highest over the prior 20 trading days. A lottery event is a MAX event for which the corresponding maximum return also ranks in the top decile across all stocks on that day. $AIDay_{it}^{SA}$ is an indicator equal to one if at least one AI-generated article/post about stock i appears on day t on SA. $AIDay_{it}^{WSB}$ is defined similarly for WSB. *HumanDay*, *Sentiment* and *Attention* measures are defined as in Table 8. All regressions include firm-characteristic controls and stock and day fixed effects. Standard errors are two-way clustered by stock and day. t-statistics are reported in parentheses. Significance levels are denoted as: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	SA		WSB		Combined	
	$MAX_{[t+1,t+5]}$	$Lottery_{[t+1,t+5]}$	$MAX_{[t+1,t+5]}$	$Lottery_{[t+1,t+5]}$	$MAX_{[t+1,t+5]}$	$Lottery_{[t+1,t+5]}$
	(1)	(2)	(3)	(4)	(5)	(6)
$AIDay^{SA}$	-0.027	-0.336*			0.020	-0.115
	(-0.37)	(-1.66)			(0.28)	(-0.54)
$HumanDay^{SA}$	-0.065	-0.365**			-0.014	-0.106
	(-1.12)	(-2.24)			(-0.25)	(-0.61)
$Attention^{SA}$	0.040***	0.148***			0.029***	0.091***
	(4.18)	(6.07)			(3.18)	(3.36)
$Sentiment^{SA}$	-0.017***	-0.010			-0.016***	-0.007
	(-3.80)	(-0.66)			(-3.61)	(-0.46)
$AIDay^{WSB}$			0.166**	0.366**	0.171**	0.403**
			(2.12)	(2.01)	(2.19)	(2.24)
$HumanDay^{WSB}$			0.002	0.139	0.017	0.190
			(0.03)	(0.98)	(0.33)	(1.39)
$Attention^{WSB}$			0.047***	0.133***	0.028***	0.079***
			(3.65)	(4.18)	(3.30)	(3.86)
$Sentiment^{WSB}$			0.009***	0.008	0.006***	0.006
			(2.69)	(0.72)	(2.70)	(0.91)
Size	-0.448***	-0.729***	-0.498***	-0.701***	-0.443***	-0.747***
	(-9.12)	(-8.78)	(-7.70)	(-6.19)	(-9.03)	(-9.13)
BM	-0.090***	-0.125***	-0.094***	-0.088	-0.086***	-0.115***
	(-4.34)	(-3.30)	(-3.68)	(-1.64)	(-4.28)	(-3.18)
IO	-0.005	-0.098	0.001	-0.091	-0.003	-0.083*
	(-0.47)	(-1.55)	(0.19)	(-1.54)	(-0.34)	(-1.76)
Analysts	-0.010	-0.040	-0.009	-0.055	-0.010	-0.045
	(-1.03)	(-1.33)	(-0.69)	(-1.12)	(-1.02)	(-1.47)
News	-0.122***	-0.154***	-0.164***	-0.175***	-0.126***	-0.157***
	(-11.02)	(-6.09)	(-11.50)	(-4.10)	(-11.53)	(-6.35)
Earnings Day	1.066***	1.543***	1.124***	1.645***	1.055***	1.502***
	(37.40)	(28.93)	(34.85)	(24.58)	(37.52)	(28.50)
$Ret_{[t-21,t-1]}$	-0.928***	-0.136***	-0.805***	-0.128***	-0.931***	-0.151***
	(-37.52)	(-7.25)	(-26.36)	(-5.03)	(-37.73)	(-8.25)
$Volatility_{[t-21,t-1]}$	-1.682***	0.008	-1.282***	0.012	-1.695***	0.009
	(-37.95)	(0.44)	(-22.79)	(0.60)	(-36.56)	(0.45)
Day FE	X	X	X	X	X	X
Stock FE	X	X	X	X	X	X
Obs.	1,424,072	680,696	635,342	274,574	1,446,735	700,185
Pseudo R^2 (%)	12.2	10.0	12.0	11.4	12.0	10.1

Appendix

A Validation of AI Detection Methods

A.1 Construction of a Ground-Truth Evaluation Sample

To evaluate the detectors, we require a corpus in which the true source of every document is known. We therefore build a test set that pairs real finance commentary with machine-generated text written to the same brief. We randomly select 1,000 long-form equity articles from Seeking Alpha and 1,000 discussion messages from r/WallStreetBets, all dated between December 2020 and November 2021. The window lies well before public LLM deployment.

We then generate AI text based on *Wall Street Journal* news. We start by collecting 2,854 articles that appeared in the Wall Street Journal "Markets," "Stocks," "U.S. Markets," and "Heard on the Street" section between December 2020 and November 2021. Each article is fed to three frontier models—GPT-4o, Claude 3.5 Sonnet, and Llama 3.3, -under two fixed instructions. Prompt 1 asks the model to recast the story in Seeking Alpha's house style, while Prompt 2 requests a Reddit r/WallStreetBets "DD" post. The exact instructions are reproduced below. The exercise yields 17,124 synthetic documents (2,854 articles times 3 models times 2 styles), bringing the evaluation set to 19,124 observations, roughly 10 percent human and 90 percent AI.

Prompt 1 - Seeking Alpha style:

Your task is to write a short-form Seeking Alpha analysis.
Here is the article I want you to write based on:

`{article}`

Important Instructions – Carefully craft a single analysis that:
Use Seeking Alpha's voice: write as a seasoned, opinionated investor, offering a clear perspective.
Paraphrase, never copy sentences from the article provided.
Total length 1,000-1,500 characters, including spaces.
Output your analysis directly, without any preamble.

Prompt 2 - Reddit style:

Your task is to write a short-form r/wallstreetbets Due Diligence (DD).
Here is the article I want you to write based on:

`{article}`

Important Instructions – Carefully craft a single analysis that:
Use r/wallstreetbets DD tone, potentially using slang and emphasizing high-risk/high-reward scenarios.
Paraphrase, never copy sentences from the article provided.

Total length 1,000-1,500 characters, including spaces.
Output your analysis directly, without any preamble.

A.2 Comparing AI Detection Methods

We compare four families of detectors that capture the main strands of the literature (See, e.g., Wu et al. (2025)).

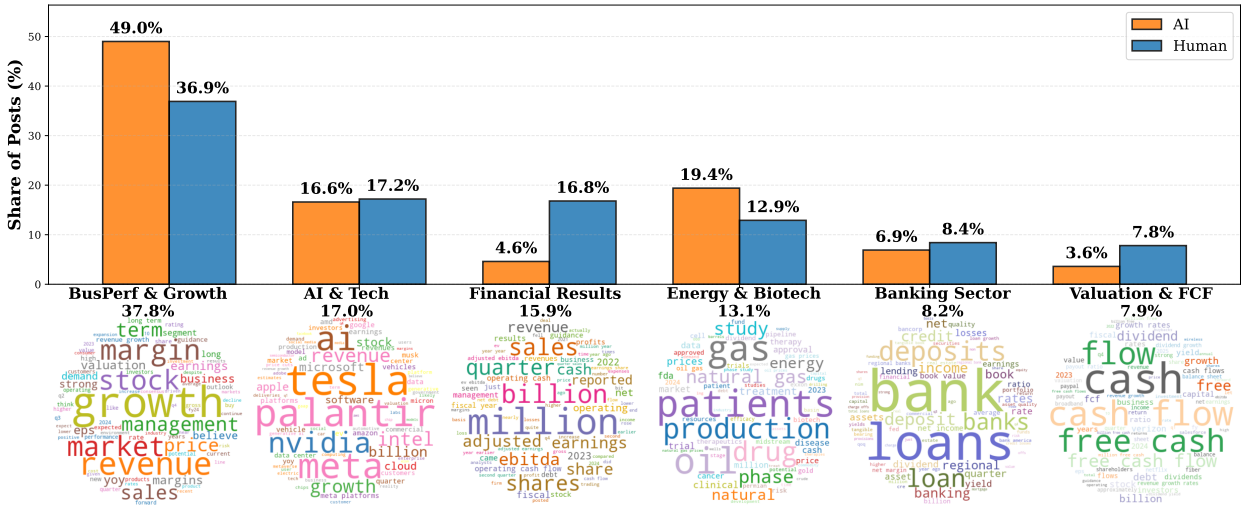
GPTZero is a proprietary ensemble of multiple lexical and probabilistic features calibrated to minimize false positives. It returns probabilities for "Human", "AI", and "MIXED", as documented in detail above. Perturbation curvature is represented by DetectGPT (Mitchell et al. 2023). The algorithm measures how log-likelihood falls when the text is lightly masked, exploiting the idea that AI prose sits at a local optimum of its own model. We implement the authors' public code and methodologies. Statistical signatures follow Su et al. (2023). We report the Normalised Perplexity Ratio (NPR) and five auxiliary scores-Log-Likelihood Ratio, entropy, token rank, raw log-likelihood, and log-rank-computed under a Llama 3.3 reference. Cut-offs are set at the 28.6-percentile of each score's distribution in the human subset, a value that equates type-I and type-II costs in a validation grid.

A.3 Detector Performance and Comparison

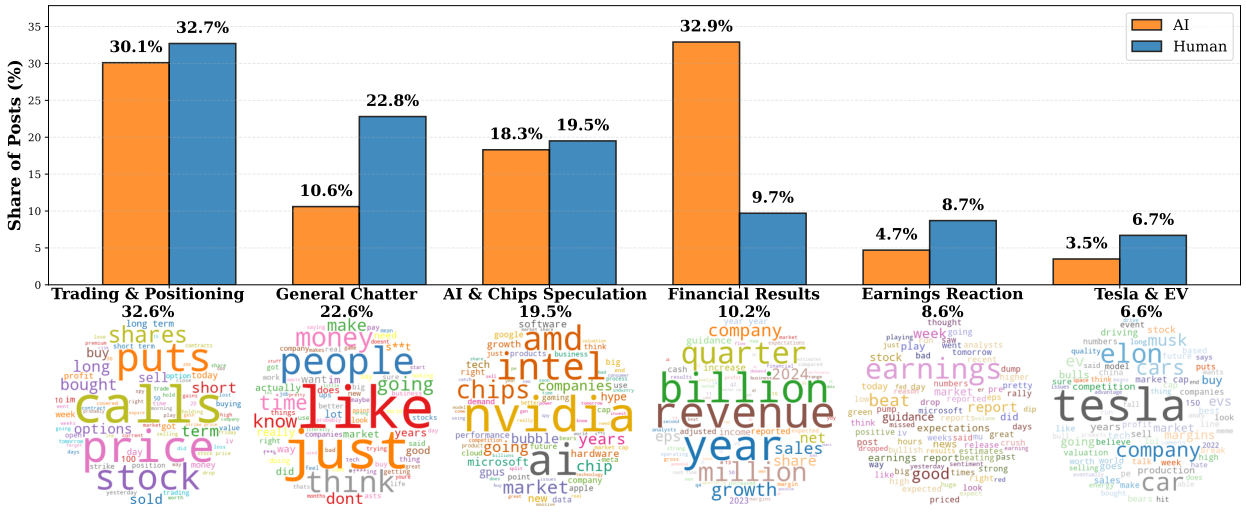
Appendix Table A2 summarises accuracy and F1 statistics averaged across the full benchmark. GPTZero identifies provenance with 98 percent accuracy and an F1 of 0.99. DetectGPT attains 81 percent accuracy ($F1 = 0.89$); NPR matches that level, and the remaining statistics hover between 76 and 80 percent. The Log-Likelihood Ratio is the strongest of the auxiliary scores, yet still trails GPTZero by roughly twenty percentage points. Because following empirical discussions rely on the accurate classification of each document, we adopt GPTZero as the detector for all subsequent analysis and treat its output as the measure of AI involvement in platform content.

A.4 Illustrative Examples of AI-Generated and Human Content

Table A3 compares human-authored content with suspected AI-generated content across both Seeking Alpha and Reddit's r/WallStreetBets. The human-written samples often exhibit a conversational, idiosyncratic tone. The suspected AI-generated texts typically display mechanical transitions, a lack of personal voice, and a rigid summary structure. For each example, we provide the excerpt context, the associated stock, the GPTZero classification probabilities, and the stylistic cues highlighted by GPTZero.



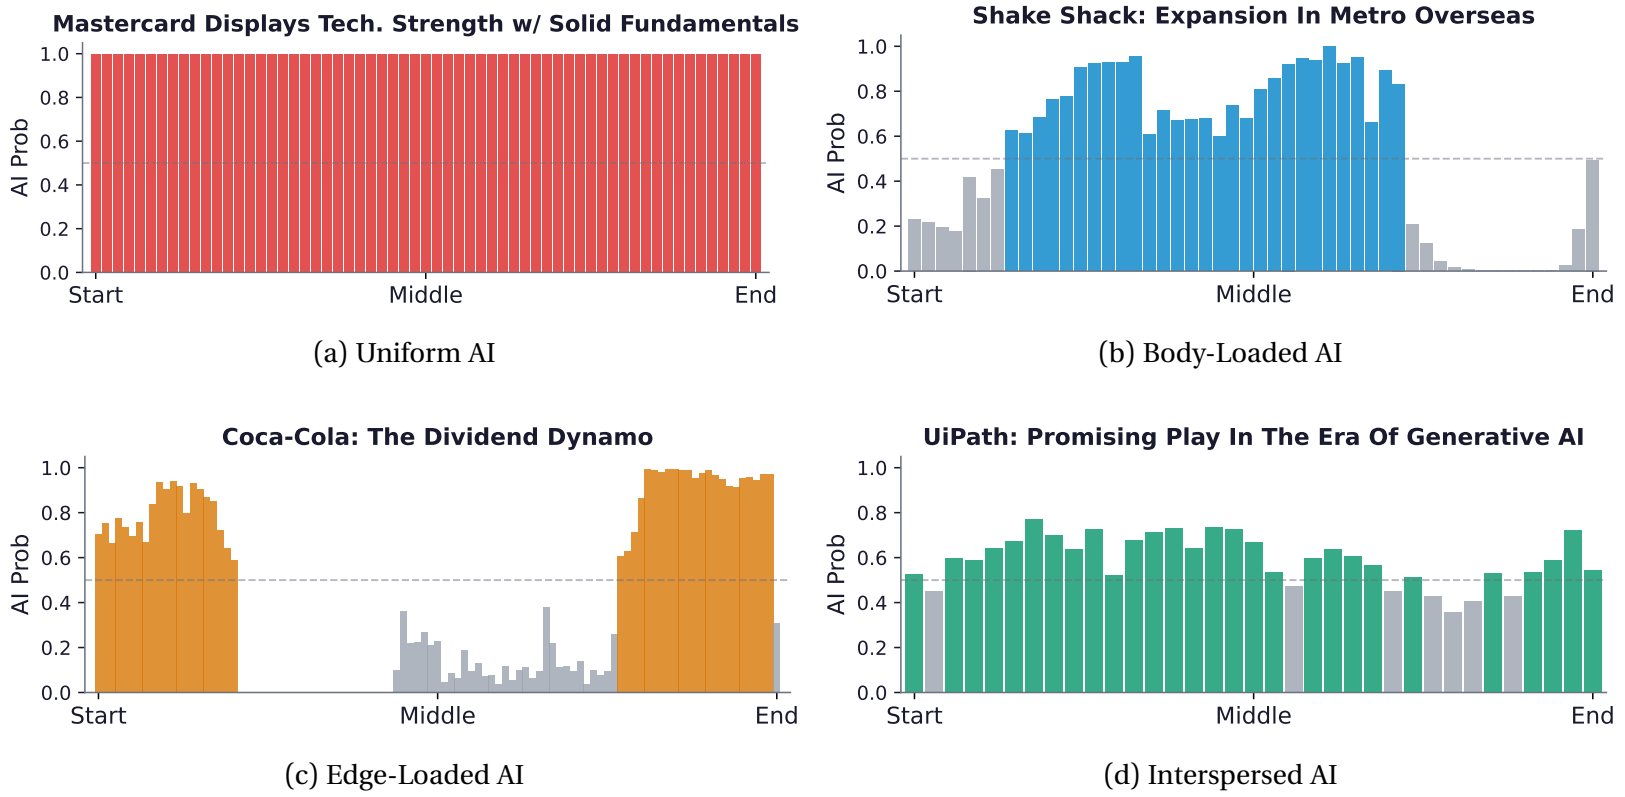
(a) Seeking Alpha



(b) WallStreetBets

Figure A1: Topic Distribution Around Earnings Announcements.

This figure displays the topic distribution for articles published in the [0, +3] window following earnings announcements. Word clouds show the key terms associated with each topic, estimated using non-negative matrix factorization (Lee and Seung, 1999) with six topics. The bar charts compare topic shares between AI-generated (orange) and human-authored (blue) content. Panel (a) shows Seeking Alpha articles (n = 7,965). Panel (b) shows WallStreetBets posts (n = 8,093).



(a) Uniform AI

(b) Body-Loaded AI

(c) Edge-Loaded AI

(d) Interspersed AI

Figure A2: Exemplar AI Usage Patterns.

This figure shows representative Seeking Alpha articles illustrating four patterns in the within-article distribution of sentences that are likely to be AI-generated. Uniform AI exhibits consistently high AI probabilities throughout the article. Body-Loaded AI concentrates AI-likely sentences in the body, with more human-written framing. Edge-Loaded AI concentrates AI-likely sentences in the introduction and conclusion. Interspersed AI alternates frequently between AI-likely and human-likely sentences. See Section 7.1 for category definitions. Each article is decomposed into ordered sentences from "Start" to "End" on the x-axis. The y-axis reports the sentence-level AI probability (AI Prob) from GPTZero. Colored bars denote AI-likely sentences (probability > 0.5), while gray bars denote human-likely sentences; the dashed line marks the 0.5 threshold.

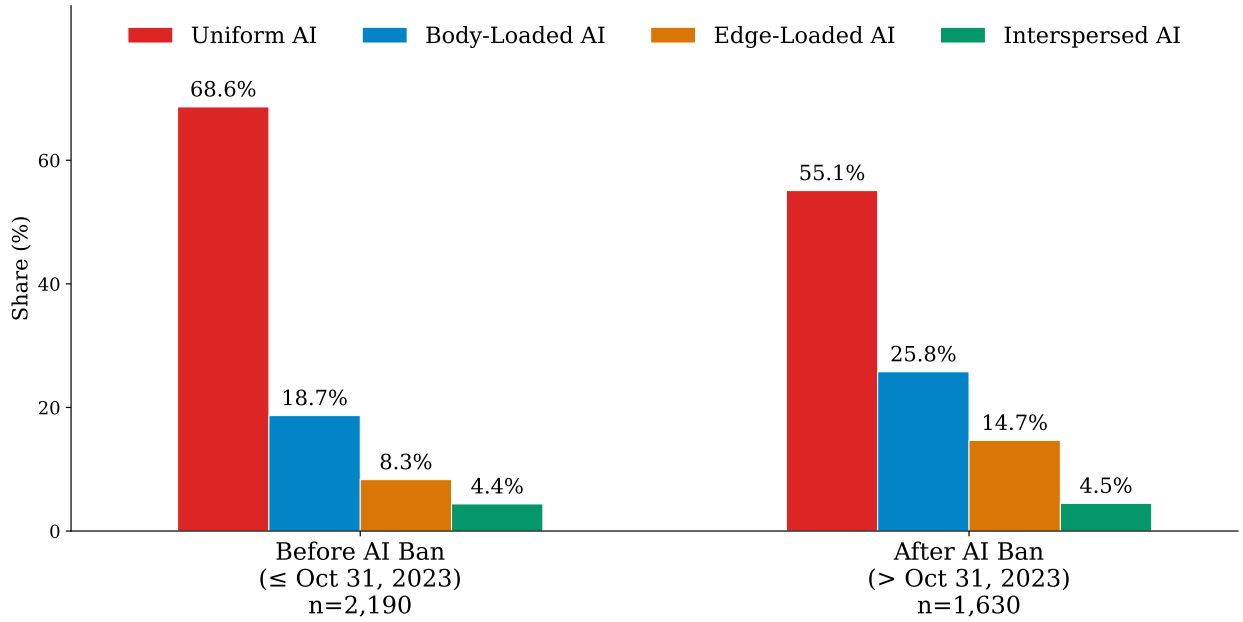


Figure A3: AI Usage Patterns Before and After the Seeking Alpha AI Ban.

This figure shows the distribution of AI usage patterns among AI articles on Seeking Alpha, comparing the pre-ban period (through October 31, 2023; n = 2,190) with the post-ban period (after October 31, 2023; n = 1,630). The sample includes articles classified as AI-Only or Mixed by GPTZero. AI usage patterns are defined based on the positional distribution of sentence-level AI probability: Uniform AI exhibits high AI probability throughout; Body-Loaded AI concentrates AI in the body with human framing; Edge-Loaded AI concentrates AI in introductions and conclusions; Interspersed AI shows frequent alternation between AI and human sentences. See Section 7.1 for category definitions.

Table A1. Description of Variables

Variable	Definition
Article-level	
AIProb_{ijt}	The sum of probability scores of content j for the AI-ONLY and MIXED categories, assigned by GPTZero.
$\text{AI-Only Prob}_{ijt}$	The probability scores of content j for the AI-ONLY categories, assigned by GPTZero.
Class_{ijt}	The classification (i.e., AI-ONLY, MIXED, HUMAN) of content j, assigned by GPTZero.
Length_{ijt}	Natural logarithm of word count of content j. For clarity, the descriptive statistics in Table 1 are based on the original, untransformed word count.
Sentiment_{ijt}	Sentiment scores of content j derived from platform-specific dictionaries (Loughran and McDonald (2011) for Seeking Alpha; Hu et al. (2025), modified for Reddit slang and emojis, for r/WallStreetBets). The sentiment score is calculated as follows: $$\text{Sentiment}_{ijt} = \frac{N_{\text{positive}} - N_{\text{negative}}}{N_{\text{positive}} + N_{\text{negative}}}$$ Where N_{positive} and N_{negative} are respectively the count of positive/negative words in the platform-specific content j. To reduce the impact of extreme values, the resulting ratio is winsorized at the 1st and 99th percentiles.
Complexity_{ijt}	Complexity scores derived from Loughran and McDonald (2024)'s dictionaries, calculated as follows: $$\text{Complexity}_{ijt} = \frac{\text{Words in Dictionary}}{\text{Total Words}}$$ Where #Words in Dictionary are the count of the specific complexity words identified in Loughran and McDonald (2024).
$\text{Quantitative}_{ijt}$	Total count of numbers mentioned divided by word count of a article, winsorized at the 1st and 99th percentiles.
Graphical_{ijt}	Total count of images divided by word count, winsorized at the 1st and 99th percentiles.
Fog Index_{ijt}	Fog Index is a readability test for English writing, calculated as follows: $$\text{Fog Index}_{ijt} = 0.4 \times \left(\frac{\text{Words}}{\text{Sentences}} + 100 \times \frac{\text{Complex Words}}{\text{Words}} \right),$$ where Words represents the total number of words in content j; Sentences represents the total number of sentences in content j; Complex Words represents the count of words with three or more syllables (excluding common suffixes like -es, -ed, -ing, proper nouns, and familiar jargon). To reduce the impact of extreme values, the resulting ratio is winsorized at the 1st and 99th percentiles.
Submissions_{ijt}	An indicator that identifies whether content j in Reddit r/WallStreetBets is a submission (top-level post).
$\text{Unfamiliarity}_{ijt}$	An indicator that identifies whether an author or user is addressing a specific firm i in their published content j for the first time within a recent six-month period.

Variable	Definition
First Time Industry$_{ijt}$	This variable is an indicator that identifies whether an author or user is addressing the specific industry of the firm i in their published content for the first time within a recent six-month period.
News$_{ijt}$	Natural logarithm of one plus the number of news articles concerning firm i mentioned in content j in the past three days ($[t-2, t]$), sourced from RavenPack.
Analysts$_{ijt}$	Natural logarithm of one plus the number of analysts covering firm i mentioned in content j, sourced from Institutional Brokers' Estimate System (IBES).
IntenseRetailBuy$_{i[t-3,t-1]}$	An indicator variable that equals 1 if firm i's retail order imbalance in the preceding three days ranks in the top 10% across-sectionally. Order imbalance is calculated as in Boehmer et al. (2021).
Comments$_{ij[t,t+10]}$	Natural logarithm of number of user comments received on an article from its publication day (t) through 10 days post-publication (t + 10).
Disagreement$_{ij[t,t+10]}$	Standard deviation of sentiment scores from user comments posted in response to content j during the period $[t, t + 10]$. The sentiment scores for these comments are calculated using the dictionary from Loughran and McDonald (2011) for Seeking Alpha content, and a dictionary from Hu et al. (2025), modified for Reddit slang and emojis, for r/WallStreetBets content.
CommentSentiment$_{ij[t,t+10]}$	Mean sentiment of comments following content j from its publication day (t) through 10 days post-publication (t + 10).
Size$_{ijt}$	Natural logarithm of the market capitalization for firm i mentioned in content j.
BM$_{ijt}$	Natural logarithm of book-to-market ratio for firm i mentioned in content j.
IO$_{ijt}$	Institutional ownership, calculated as the percentage of common shares outstanding of firm i mentioned in content j, held by institutional investors as of the last quarter ending before day t.
Author Articles$_{ijt}$	Natural logarithm of one plus the number of articles published by a given author in the past six months.
Author AI Articles$_{ijt}$	Natural logarithm of one plus the number of AI-ONLY or MIXED articles by a given author in the past six months.
Author Comments$_{ijt}$	Natural logarithm of one plus the average number of comments per article within 21 days, where the average is taken across all articles that a given author publishes in the past six months.
Author level	
Author Article Count$_a$	Number of articles that author/user a published during the sample period.
# Covered Stocks$_a$	Number of distinct stocks that author/user a covered at least once during the sample period.
# Covered Industries$_a$	Number of distinct CRSP four-digit SIC industries across the set of tickers covered by author/user a during the sample period.
Stock-day level	
AIDay$_{i,t}$	An indicator that equals one if at least one article about firm i published on day t is classified as AI-ONLY or MIXED by GPTZero.
Article$_{i,t}$	The number of stock-specific articles about firm i published on day t across the two platforms.

Variable	Definition
$\text{AVOL}_{i[t+1,t+5]}$	Average abnormal volume for firm i over the 5 trading days from $t+1$ to $t+5$: $$\text{AVOL}_{i[t+1,t+5]} = \frac{1}{5}\sum_{s=1}^{5}\log(1+\textit{Volume}_{i,t+s}) - \frac{1}{31}\sum_{s=-41}^{-11}\log(1+\textit{Volume}_{i,t+s}),$$ where *Volume* is the dollar volume, calculated as the numbers of shares traded multiplied by the price per share.
$\text{Volatility}_{i[t+1,t+5]}$	Standard deviation of the daily stock returns for firm i over the 5 trading days from $t+1$ to $t+5$.
$\text{Spread}_{i[t+1,t+5]}$	Average daily volume-weighted effective bid-ask spread for firm i over the 5 trading days from $t+1$ to $t+5$.
$\text{MAX}_{i[t+1,t+5]}$	A MAX event is when the daily return is the highest over the prior 20 trading days.
$\text{Lottery}_{i[t+1,t+5]}$	A lottery event is a MAX event where the corresponding 20-day maximum return also ranks in the top decile compared to all other stocks on that same day.
$\text{CAR}_{i[t+1,t+5]}$	Cumulative buy-and-hold market-adjusted returns over the 5 trading days from $t+1$ to $t+5$.
OIB_{it}	Retail order imbalance for stock i on day t, computed as $$OIB_{i,t} = \frac{\#\text{retail buy trades}_{i,t} - \#\text{retail sell trades}_{i,t}}{\#\text{retail buy trades}_{i,t} + \#\text{retail sell trades}_{i,t}}$$ using marketable retail trades identified following Boehmer et al. (2021).
Sentiment_{it}	The average sentiment score of articles about firm i on day t. Value set to a neutral score of zero for days without any content of firm i, following Cookson et al. (2024)
Attention_{it}	Natural logarithm of one plus the number of articles on SA (or posts on WSB) of firm i on day t.
Size_{it}	Natural logarithm of the market capitalization for firm i.
BM_{it}	Natural logarithm of book-to-market ratio for firm i.
IO_{it}	Institutional ownership, calculated as the percentage of common shares outstanding of firm i held by institutional investors as of the last quarter ending before day t.
Analysts_{it}	Natural logarithm of one plus the number of analysts covering firm i, sourced from Institutional Brokers' Estimate System (IBES).
News_{it}	Natural logarithm of one plus the number of news articles concerning firm i in the past 21 trading days, sourced from RavenPack.
$\text{Ret}_{i[t-21,t-1]}$	Cumulative returns over the preceding 21 trading days from $t-21$ to $t-1$.
$\text{Volatility}_{i[t-21,t-1]}$	Standard deviation of the daily stock returns for firm i over the preceding 21 trading days from $t-21$ to $t-1$.
$\text{Spread}_{i[t-21,t-1]}$	Average daily volume-weighted effective bid-ask spread for firm i over the preceding 21 trading days from $t-21$ to $t-1$.

Table A2. AI Detection Method Comparison

This table compares AI text detection methods on our benchmark dataset, reporting overall accuracy and F1 scores. Accuracy is the proportion of documents whose provenance (AI-generated or human-written) is correctly identified, and the F1 score is the harmonic mean of precision and recall, providing a balanced assessment of classification performance. Higher values indicate better performance on both metrics. The comparison includes GPTZero, DetectGPT, NPR, Log-Likelihood Ratio (LLR), Entropy, Rank, Log Likelihood, and Log Rank.

Detection Method	**Accuracy**	**F1 Score**
GPTZero	98%	99%
DetectGPT	81%	89%
NPR	80%	89%
LLR	78%	87%
Entropy	76%	86%
Rank	76%	86%
Log Likelihood	76%	86%
Log Rank	76%	86%

Table A3. **Examples of SA Articles, with GPTZero Output**

SA Article 1

Title: Nike Is Running Well, But Its Valuation Is A Hurdle
Ticker: NKE (Nike Inc.)
Date: Feb. 10, 2024
User: Khen Elazar
Analyst Recommendation: HOLD
Link: https://seekingalpha.com/article/4669281-nike-is-running-well-but-its-valuation-is-a-hurdle

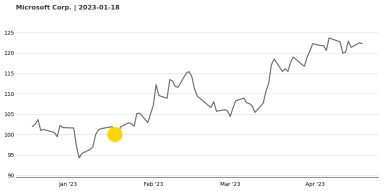


Excerpt: These risks, while manageable, necessitate strategic planning and operational flexibility to mitigate potential impacts on revenue and profitability. Given the current market valuation, the shares are overvalued. Taking into account NIKE's solid foundation, growth prospects, and the outlined risks, the verdict leans towards a HOLD.

GPTZero Detection: AI-generated
AI-Only (97.75%), Mixed (0.87%), Human (1.37%)
Explanation: Mechanical Transitions, Task-Oriented, Formulaic, Lacks Creativity

SA Article 2

Title: Microsoft Doubling Down On OpenAI
Ticker: MSFT (Microsoft Corp.)
Date: Jan. 18, 2023
User: Ivy Global Insights
Analyst Rating: BUY
Link: https://seekingalpha.com/article/4570487-microsoft-stock-doubling-down-openai



Excerpt: Everyone knows ChatGPT. It is a smart coder, an intelligent writer, and an elegant Chat Bot. We all get used to search engines that pops up random pages with key words, and we dig into those links on the first or second page of search results in hoping to locate what we want. In contrast, ChatGPT now emerges as powered by all wisdom in the universe and it can actually hear us and respond to us. For the first time, search becomes a peaceful and delightful human-like dialogue.

GPTZero Detection: Human-written
Human (99.25%), Mixed (0.20%), AI-Only (0.54%)



Table A4. **Examples of WSB Articles, with GPTZero Output**

WSB Article 1

Title: 40K+ in Intel INTC ! I love USA chip PART 2
Ticker: INTC (Intel Inc.)
Date: December, 10, 2024
User: /Intelligent-Cellist6
Link: https://www.reddit.com/r/wallstreetbets/comments/1hbakxw/40k_in_intel_intc_i_love_usa_chip_part_2/

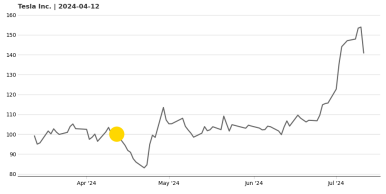


Excerpt: Going **all in** on $INTC!
If you've seen my prior post, you know I crushed it with LEAPS, riding Intel from **$18** to **$26**. Now, I'm back for round two. Currently sitting on **2K shares** and planning to **load up another 3K with margin**-yes, I'm that bullish.
Here's why I'm doubling down...
Undervalued gem - At these levels, $INTC is a straight-up steal.
I see huge upside potential here-like staring into the abyss of tendies. Who's riding with me?

GPTZero Detection: AI-generated
AI-Only (99.98%), Mixed (0.02%), Human (0%)
Explanation: Promotional Tone, Artificial Simplicity, Speculative Focus

WSB Article 2

Title: TSLA earnings play
Ticker: TSLA (Tesla Inc.)
Date: Apr. 12, 2024
User: u/Fragrant_Swing9987
Link: https://www.reddit.com/r/wallstreetbets/comments/1c2fztr/tsla_earnings_play/



Excerpt: Hello regards. Today I thought of a potential TSLA earnings call play. Before you guys bomb me with hate notice that the earning expectations are very low for TSLA, with a EPS estimate of 0.36 (compared to 0.6 Q4 2023). Thus, I think a potential outperformance can occur at TSLA earnings shooting up the price. What are your thoughts ? (Not financial advice, my brain is smooth)

GPTZero Detection: human-written
Human (98.71%), Mixed (0%), AI-Only (1.28%)



Table A5. Textual Attributes of AI-Generated Articles

This table examines the association between GPTZero's AI probability score and textual attributes of articles/posts at the article level using the following regression:

$$AIProb_{ijt} = \gamma' \boldsymbol{X}_{ijt} + \epsilon_{ijt}$$

where $AIProb_{ijt}$ (in percentage points) is the AI probability score for content j about firm i, posted on day t. $\boldsymbol{X}_{ijt}$ is content j's characteristics (standardized to zero mean and unit variance). Columns 1–2 report results for articles published on Seeking Alpha; Columns 3–4 report results for posts on Reddit's r/WallStreetBets. Standard errors are two-way clustered by stock and day; t-statistics are in parentheses. $^{*}p$ <.1; $^{**}p$ <.05; $^{***}p$ <.01.

	SA		WSB			
	(1)	(2)	(3)	(4)	(5)	(6)
Length	0.345***	-0.043	2.737***	1.479***	2.106***	1.286***
	(3.44)	(-0.32)	(16.88)	(10.13)	(13.70)	(8.49)
Sentiment	2.943***	0.774***	0.858***	0.220***	0.804***	0.216***
	(22.31)	(8.02)	(8.82)	(3.15)	(8.83)	(3.07)
Complexity	1.866***	0.410***	0.577***	0.260***	0.550***	0.255***
	(14.94)	(5.04)	(6.51)	(2.88)	(6.55)	(2.83)
Fog Index	1.771***	3.839***	0.397***	0.171**	0.424***	0.178**
	(13.29)	(20.35)	(6.97)	(2.35)	(7.66)	(2.44)
Graphical	-0.402***	-0.962***	0.054	-0.030	-0.382***	-0.164
	(-3.27)	(-7.35)	(0.63)	(-0.24)	(-4.10)	(-1.37)
Quantitative	-1.332***	-0.313***	0.278***	0.200*	0.139*	0.170
	(-9.78)	(-2.87)	(3.17)	(1.81)	(1.75)	(1.58)
Submissions					6.964***	2.834***
					(9.00)	(4.09)
Day FE	X	X	X	X	X	X
Stock FE	X	X	X	X	X	X
Author FE		X		X		X
Obs.	57,581	57,581	64,848	64,848	64,848	64,848
Adj. R^2 (%)	9.9	59.5	10.3	56.9	12.0	57.1

Table A6. Rhetorical Analysis: Classification Examples

This table presents illustrative examples of rhetorical classification from both platforms, including the classifier's reasoning. The logos-dominant and ethos-dominant examples are AI-ONLY and HUMAN articles on SA detected by GPTZero. The pathos-dominant example is AI-ONLY article on WSB detected by GPTZero. Each example shows the article headline, key evidence sentence, composition scores, effectiveness ratings, and the model's explanation.

Example: Logos-Dominant

Headline: "Amazon's Stock May See A Sharp Decline Following Results" [SA]
Composition: Logos 88% / Pathos 4% / Ethos 8%
Effectiveness: Logos 9 / Pathos 3 / Ethos 6
Key Evidence: "Analysts predict a 13.5% decline in earnings to $0.46 per share compared to last year, while revenue is forecasted to increase by 7.1% to $124.7 billion."
LLM Reasoning: The article is overwhelmingly data-driven and analytical (earnings forecasts, valuation multiples, option pricing, technical patterns), with only light emotional language and modest implied credibility from detailed trade mechanics, making LOGOS both dominant and highly effective.

Example: Pathos-Dominant

Headline: "What Are Your Moves Tomorrow, November 14, 2024" [WSB]
Composition: Logos 5% / Pathos 90% / Ethos 5%
Effectiveness: Logos 2 / Pathos 9 / Ethos 3
Key Evidence: "Imagine selling RKLB at $8, then watching-helpless-as it rockets to $22. This isn't just a miss. This is portfolio-crushing, FOMO-fueled regret."
LLM Reasoning: The text is almost entirely an emotional FOMO narrative with vivid regret and urgency, minimal factual or logical content, and only weak implied credibility from shared trader experience.

Example: Ethos-Dominant

Headline: "Sentiment Speaks: $2,428 Target For Gold Now In Focus"[SA]
Composition: Logos 35% / Pathos 25% / Ethos 40%
Effectiveness: Logos 7 / Pathos 6 / Ethos 9
Key Evidence: "And, those that have tracked us over that time have recognized the accuracy of not only our short term calls but of our long term ones as well."
LLM Reasoning: The article is primarily a catalogue of past market calls and trade outcomes to showcase the author's track record (strong ethos), supported by some specific price targets and projections (moderately effective logos) and framed with language about contrarian courage and market crowd behavior that adds an emotional, but secondary, pathos appeal.

Table A7. AI Persuasion Effectiveness on WallStreetBets, by Pathos Tercile

This table compares rhetorical effectiveness betweenm AI-generated and human-authored content on WallStreetBets. Posts are sorted into terciles based on their *pathos* share (the percentage of text classified to emotional appeals by our rhetorical analyzer). For each tercile, we report mean effectiveness ratings (on a 0–10 scale) separately for AI-generated posts (defined as those classified by GPTZero as AI-ONLY or MIXED) and human-authored posts. The "AI-Human" columns report the difference in means, with significance determined by two-sample t-tests. *p <.1; **p <.05; ***p <.01.

	Pathos Effectiveness			Logos Effectiveness			Ethos Effectiveness		
Tercile	AI	Human	AI-Human	AI	Human	AI-Human	AI	Human	AI-Human
High	7.37	7.08	+0.28***	2.55	2.60	−0.06	3.20	3.09	+0.11
Medium	7.23	6.57	+0.66***	4.86	3.98	+0.89***	4.07	3.72	+0.35***
Low	4.48	4.89	−0.41***	6.60	5.64	+0.96***	4.48	4.11	+0.37***

Table A8. Analysis of the Seeking Alpha AI Ban

This table replicates key results before and after the Seeking Alpha AI ban (October 31, 2023). We interact key variables with Pre-Ban and Post-Ban indicators. Columns 1–6 corresponds to Table 3–7. Columns 7–9 corresponds to Table 8. Columns 10–11 corresponds to Table 9. All specifications include day and stock fixed effects; columns 1–3 additionally include author fixed effects. Standard errors are two-way clustered by stock and day. t-statistics in parentheses. $^{*}p < .10$; $^{**}p < .05$; $^{***}p < .01$.

	AIProb			Commt. Sent.	CAR[1, 5]		Volume	Volatility	Spread	MAX	Lottery
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Unfamiliarity × Pre-Ban	1.107***										
	(4.77)										
Unfamiliarity × Post-Ban	0.024										
	(0.10)										
News × Pre-Ban		−0.262*									
		(−1.94)									
News × Post-Ban		−0.157									
		(−1.24)									
Analysts × Pre-Ban		−0.313*									
		(−1.66)									
Analysts × Post-Ban		−0.346*									
		(−1.82)									
Intense Rtl. Buy × Pre-Ban			0.600								
			(0.97)								
Intense Rtl. Buy × Post-Ban			−0.695								
			(−1.15)								
AIProb × Atcl. Sent. × Pre-Ban				−0.092	0.100						
				(−1.21)	(0.59)						
AIProb × Atcl. Sent. × Post-Ban				0.031	0.701**						
				(0.37)	(2.24)						
AIDay × OIB × Pre-Ban						0.294					
						(1.34)					
AIDay × OIB × Post-Ban						0.613					
						(1.53)					
AIDay × Pre-Ban						−0.611*	−0.031	−0.006	−0.041***	0.020	−0.249
						(−1.70)	(−1.17)	(−0.24)	(−5.94)	(0.24)	(−1.04)
AIDay × Post-Ban						−0.518	−0.006	−0.033	−0.030***	−0.097	−0.453*
						(−1.22)	(−0.20)	(−1.26)	(−3.84)	(−1.15)	(−1.76)
Controls	X	X	X	X	X	X	X	X	X	X	X
Day FE	X	X	X	X	X	X	X	X	X	X	X
Stock FE	X	X	X	X	X	X	X	X	X	X	X
Author FE	X	X	X								
Obs.	54,429	54,429	54,429	54,429	54,070	1,370,181	1,424,099	1,424,099	1,358,845	1,424,072	680,696
Adj. R^2/Pseudo R^2 (%)	59.6	59.5	60.0	19.1	8.6	3.2	17.7	38.9	82.4	12.2	10.0

Table A9. AI-Generated Content and User Engagement

This table examines the association between AI probability and user engagement using the following regression:

$$Engagement_{ij[t,t+10]} = \beta_1 AIProb_{ijt} + Controls + FE + \epsilon_{ijt}.$$

$Engagement_{ij[t,t+10]}$ captures user responses to content j about firm i over the window $[t, t+10]$. It is measured as $\log(1 + Comments)$, the natural logarithm of one plus the number of comments received in columns 1–2, and *Disagreement*, the standard deviation of comment sentiment scores (requring at least two comments; standardized to zero mean and unit variance) in columns 3–6. $AIProb_{ijt}$ is the estimated probability, ranging from 0 to 1, that content j covering firm i on day t is AI-generated. *Controls* include article characteristics, firm characteristics, and measures of the author's prior activity and performance (standardized to zero mean and unit variance). All specifications include stock, day, and author fixed effects. Standard errors are two-way clustered by stock and day; t-statistics are in parentheses. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

	Comments		Disagreement	
	SA	WSB	SA	WSB
	(1)	(2)	(3)	(4)
AIProb	-0.065**	-0.014	-0.066**	0.105
	(-2.48)	(-0.26)	(-1.99)	(0.47)
Length	0.114***	0.088***	0.039***	0.124***
	(11.71)	(12.20)	(4.10)	(3.38)
Sentiment	-0.080***	-0.020***	-0.002	0.008
	(-11.74)	(-5.16)	(-0.36)	(0.38)
Complexity	-0.032***	-0.004	-0.003	-0.001
	(-5.15)	(-0.75)	(-0.62)	(-0.04)
Fog Index	-0.033***	-0.014**	0.024**	-0.052**
	(-3.91)	(-2.54)	(2.57)	(-2.12)
Graphical	-0.012	0.005	0.011	0.069
	(-1.65)	(0.80)	(1.37)	(1.61)
Quantitative	-0.042***	0.032***	0.018**	0.019
	(-6.33)	(7.90)	(2.34)	(0.70)
Size	0.199**	0.041	0.155**	-0.028
	(2.44)	(0.50)	(2.29)	(-0.06)
BM	-0.001	-0.006	-0.011	0.039
	(-0.06)	(-0.28)	(-0.49)	(0.37)
News	0.072***	0.013	0.013*	-0.016
	(9.52)	(1.26)	(1.71)	(-0.28)
Analysts	0.007	0.025*	-0.006	0.048
	(0.67)	(1.80)	(-0.42)	(0.53)
IO	-0.056**	0.008	0.021	-0.023
	(-2.01)	(0.33)	(0.62)	(-0.19)
Author Article	-0.015**	0.014**	-0.007	0.040
	(-2.06)	(2.00)	(-0.71)	(1.40)
Author AI Article	0.006	0.006	0.010	-0.039
	(1.25)	(0.99)	(1.55)	(-1.14)
Author Comments	0.012*	-0.026***	0.006	-0.042
	(1.84)	(-4.89)	(0.60)	(-1.33)
		(53.81)		(6.44)
Day FE	X	X	X	X
Stock FE	X	X	X	X
Author FE	X	X	X	X
Obs.	54,429	63,400	38,376	20,883
Adj. R^2 (%)	64.2	43.1	10.0	18.4

Table A10. AI Days and Market Outcomes: Controlling for Earnings Announcements

This table replicates the market outcome regressions under two alternative treatments of earnings announcements: Panel A excludes stock-day observations that fall within the $[-1,+1]$ window around an earnings announcement, and Panel B additionally controls for upcoming earnings announcements. The dependent variables are cumulative abnormal returns over $[t+1, t+5]$ (column 1), abnormal trading volume (column 2), return volatility (column 3), effective bid-ask spread (column 4), maximum daily return indicator (column 5), and lottery-like return indicator (column 6). In the specification controlling for upcoming earnings, $I^{Earnings}_{[t+1,t+5]}$ is an indicator equal to one if stock i has an earnings announcement scheduled within the subsequent five trading days. $AIDay^{SA}it$ ($AIDay^{WSB}it$) is an indicator equal to one if at least one AI-generated article about stock i appears on Seeking Alpha (WallStreetBets) on day t. $\text{OIB}_{i,t}$ is the net retail order imbalance for stock i on day t. All regressions include controls for firm size, book-to-market, institutional ownership, news coverage, analyst coverage, past 21-day returns and volatility, platform attention and sentiment, and an earnings-day indicator. Stock and day fixed effects are included. Standard errors are two-way clustered by stock and day. t-statistics are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	A. Excluding Earnings-Announcement Windows					
	$CAR_{[t+1,t+5]}$ (1)	Volume (2)	Volatility (3)	Spread (4)	MAX (5)	Lottery (6)
$AIDay^{SA}$	−0.903**	0.027	0.012	−0.017***	0.076	0.075
	(−2.37)	(1.23)	(0.58)	(−5.30)	(1.01)	(0.29)
$AIDay^{WSB}$	0.047	0.079**	0.066	0.007	0.166**	0.416**
	(0.13)	(2.44)	(1.48)	(1.13)	(2.00)	(2.05)
$AIDay^{SA} \times OIB$	0.490**					
	(2.19)					
$AIDay^{WSB} \times OIB$	0.650					
	(1.64)					
Controls	X	X	X	X	X	X
Day FE	X	X	X	X	X	X
Stock FE	X	X	X	X	X	X
Obs.	1,325,524	1,379,723	1,379,723	1,254,599	1,379,697	666,357
Adj. R^2/Pseudo R^2 (%)	3.2	17.5	39.4	88.6	11.9	10.5
	B. Controlling for Upcoming Earnings Announcements					
	$CAR_{[t+1,t+5]}$ (1)	Volume (2)	Volatility (3)	Spread (4)	MAX (5)	Lottery (6)
$AIDay^{SA}$	−0.573	0.026	0.030	−0.016***	0.031	−0.105
	(−1.56)	(1.20)	(1.54)	(−5.45)	(0.44)	(−0.49)
$AIDay^{WSB}$	−0.085	0.090***	0.077*	0.004	0.152*	0.400**
	(−0.23)	(2.86)	(1.82)	(0.85)	(1.89)	(2.16)
$AIDay^{SA} \times OIB$	0.435**					
	(2.06)					
$AIDay^{WSB} \times OIB$	0.477					
	(1.06)					
$I^{Earnings}_{[t+1,t+5]}$	−0.116	0.521***	0.227***	0.002	1.208***	2.035***
	(−0.97)	(39.11)	(16.62)	(0.75)	(38.13)	(31.30)
Controls	X	X	X	X	X	X
Day FE	X	X	X	X	X	X
Stock FE	X	X	X	X	X	X
Obs.	1,391,849	1,446,761	1,446,761	1,365,371	1,446,735	700,185
Adj. R^2/Pseudo R^2 (%)	3.2	20.1	41.4	88.6	12.8	12.8

Internet Appendix

Contents

A Construction of Seeking Alpha Author Characteristics **2**

B Rhetorical Classification Prompt **2**

C AI Usage Analysis **4**
C.1 Overview . . . 4
C.2 Constructing the Empirical PDF . . . 5
C.3 Benchmark PDF Definitions . . . 5
C.4 Distance Metric and Transition Rate . . . 6
C.5 Classification Algorithm . . . 7
C.6 Output Variables . . . 7
C.7 Limitations . . . 7

List of Appendix Tables

IA1 AI-Generated Content and Sentiment Alignment: Robustness . . . 8
IA2 WallStreetBets Analysis: Alternative Samples . . . 9

A Construction of Seeking Alpha Author Characteristics

We construct measures of Seeking Alpha author characteristics from their public biography pages. We preprocess biography text by removing HTML tags, normalizing whitespace, and standardizing punctuation to obtain a cleaned biography string. We define *No Bio* to equal one for authors with missing or empty biography fields.

We define *Certified* to equal one if an author lists at least one professional certification, including CFA (Chartered Financial Analyst), CPA (Certified Public Accountant), FRM (Financial Risk Manager), CFP (Certified Financial Planner). We identify advanced educational credentials when an author lists MBA or PhD degrees.

We classify professional roles based on self-descriptions: portfolio manager ("portfolio manager" or "fund manager"), trader ("trader" or "trading"), and advisors ("advisor," "adviser," or "consultant"). Because "analyst" is widely used on Seeking Alpha as a broad label instead of a specific financial occupation, we classify an author as an analyst only when "analyst" appears in a clearly financial context, such as being modified by financial descriptors (e.g., "equity," "credit," "macro," "quant," "hedge fund," and close variants), and we explicitly exclude non-financial uses (e.g., "systems analyst," "web analyst," "marketing analyst").

We also infer whether an author is likely a native English speaker based on the displayed names. We preprocess display names to remove non-name tokens, including professional suffixes (e.g., "CFA," "PhD," "MD"), legal-entity suffixes (e.g., "LLC," "Inc," "Ltd"), and honorifics (e.g., "Mr," "Dr"). We then exclude likely institutional accounts using corporate identifiers (e.g., "Capital," "Fund," "Research," "Global," "Media") and non-personal patterns (e.g., special characters such as "&" or "@"). For the remaining names, we using Namsor, a name-classification API (https://namsor.app/), to obtain a predicted country of origin. We define an English-language indicator equal to one when the predicted origin corresponds to an English-primary official-language country (the United States, United Kingdom, Canada, Ireland, Australia, and New Zealand), and zero otherwise. Prior studies use Namsor to infer gender, ethnicity and geographic origin (Lu, Naik, and Teo 2024; Condie et al. 2023). This name-based measure should be interpreted as a proxy for English-primary origin rather than a direct measure of native-speaker status. It is therefore noisy, and the associated measurement error likely attenuates the estimated coefficient toward zero.

B Rhetorical Classification Prompt

This section describes the prompts used to classify rhetorical strategies in contents on Seeking Alpha and WallStreetBets. We use GPT-5.1 with temperature set to 0 to ensure reproducibility. The prompt is designed to be platform-neutral, avoiding any mention of specific platforms that might bias the classifier. Below, we report the complete system prompt.

Rhetorical Classification System Prompt

```
You are a rhetorical analyst for financial social-media research.
You will be given a text about a financial asset. The source could be a formal
research article or a social media discussion.
Your task is to analyze the text based on the Aristotelian Triad (Logos, Pathos,
Ethos). You must evaluate the text on two distinct dimensions:
1. Composition: What percentage of the text is dedicated to each mode? (Must sum
to 100%)
2. Effectiveness: How convincing/persuasive is each mode within this specific text?
(Score 0-10)

DEFINITIONS & INDICATORS

(1) LOGOS - The Appeal to Logic & Fact
```

Definition: Persuasion through reasoning, evidence, financial mechanics, and objective exposition.
NOTE: Plain, factual description (e.g., "The company was founded in 1990") counts as LOGOS because it provides the evidentiary basis for the argument.
Indicators:
- Structured argumentation (premise → evidence → conclusion)
- Quantitative analysis: citing valuation models, financial ratios, or price targets
- Conditional logic (e.g., "if yields rise, then X happens")
- References to primary documentation (earnings reports, filings, historical data)
- Tone: Objective, analytical, or neutral

(2) PATHOS – The Appeal to Emotion

Definition: Persuasion through psychological triggers, urgency, and shared identity.
Indicators:
- Emotional vocabulary (excitement, fear, outrage, hope)
- Imagery and metaphors (visual language, analogies)
- Urgency or behavioral triggers (calls to action, FOMO, warnings)
- In-group terminology: Use of community-specific slang, memes, or tribal identifiers
- Tone: Hyperbolic, intense, or highly subjective

(3) ETHOS – The Appeal to Credibility

Definition: Persuasion based on the author's character, authority, or risk profile.
Indicators:
- Expertise: Mention of professional credentials, deep historical experience, or track record
- "Skin in the Game": Disclosure of personal financial stake, specific positions held, or personal risk exposure
- Transparency: Clear articulation of risks, admission of uncertainty/losses, or balanced disclaimers
- Reputation: Appeals to the author's standing within their specific community or references to authoritative third parties

SCORING GUIDELINES

A. Composition Scores (Sum must = 100%)
- Estimate the "share of voice" or volume of text dedicated to each mode
- Rule for "Neutral" Text: Assign purely informational or descriptive text to LOGOS, as it serves as the factual foundation of the argument
- Note: A single sentence can contain multiple modes, but prioritize the primary intent of the sentence

B. Effectiveness Scores (0 - 10)
- Rate the quality and persuasive impact of the mode as used in the text
- 0: Absent or Counter-productive (e.g., logical fallacies, obviously fake emotion)
- 5: Competent/Average (e.g., standard arguments, typical expressions of sentiment)
- 10: Highly Potent (e.g., irrefutable data analysis, emotionally contagious

```
  narrative, or undeniable proof of authority)

OUTPUT FORMAT

Return a strict JSON object with the following structure:
{
  "logos_percent": [Integer 0-100],
  "pathos_percent": [Integer 0-100],
  "ethos_percent": [Integer 0-100],
  "logos_effective_rating": [Integer 0-10],
  "pathos_effective_rating": [Integer 0-10],
  "ethos_effective_rating": [Integer 0-10],
  "dominant_style": "LOGOS", "PATHOS", "ETHOS", or "BALANCED",
  "key_evidence_sentence": "Verbatim sentence from text...",
  "reasoning": "One concise sentence explaining the scoring logic."
}
```

C AI Usage Analysis

This appendix describes our methodology for classifying *how* AI is used in financial content.

C.1 Overview

AI can play multiple roles in content generation, including idea generation, background research, narrative framing, elaboration or polishing of human-written text, or near-complete substitution for human effort. Although we do not observe the production workflows of an author directly, the placement of AI-generated text within an article (e.g., the beginning, the body, the end, or throughout) provides indirect but informative evidence on how human and AI inputs are combined.

Different modes of AI use naturally imply different integration points within the text. For example, AI employed for broad drafting or rewriting should appear throughout an article, whereas AI used more selectively is likely to concentrate in particular sections. Motivated by this intuition, we use sentence-level AI probability scores from GPTZero to classify articles based on the positional distribution and clustering of AI-like sentences. This approach allows us to distinguish between episodic, assistive uses of AI and more pervasive forms of substitution.

We define four AI usage categories. First, *Uniform AI* captures articles in which AI probability is high and relatively flat across the full text, with little systematic variation by section; this pattern is consistent with pervasive AI involvement, either because the author relies on AI for most content generation or because AI is used to rewrite an initial human sketch or draft across the entire article.

Second, *Body-Loaded AI* describes articles in which AI probability is comparatively low near the beginning and end. This pattern suggests that authors selectively deploy AI for substantive discussion, such as background information, quantitative elaboration, or supporting analysis, while retaining greater human involvement in framing. Such usage is consistent with AI acting as a complement to human analysis by expanding or clarifying informational content.

Third, *Edge-Loaded AI* applies when AI probability is elevated at the start and/or end of the article, indicating that AI-generated sentences are concentrated in framing sections, such as crafting introductions or conclusions, while relying more heavily on human input for core analytical content.

Finally, *Interspersed AI* captures articles in which AI-like sentences are alternate frequently with human-like sentences throughout the text, with no strong concentration in any single section. This pattern is consistent

with iterative human-AI revision at a granular, perhaps sentence-by-sentence level.

To implement the classification, we first construct an empirical probability density function (PDF) for each article that represents the positional distribution of AI probability across sentences (Section C.2). Second, we define benchmark PDFs corresponding to each positional usage pattern (Section C.3). Third, we compute L1 distances between each article's empirical PDF and the benchmarks, along with a transition rate to detect interspersed editing (Section C.4). Finally, we describe the classification algorithm that assigns articles to categories based on these metrics (Section C.5).

C.2 Constructing the Empirical PDF

Our classification framework is based on the *positional distribution* of AI probability within articles. Our approach treats the normalized AI probabilities across sentence positions as a probability density function (PDF) and measures the distance between each article's empirical PDF and a set of benchmark PDFs representing archetypal usage patterns.

For each article with n sentences, let $a_i \in [0,1]$ denote the AI-generated probability for sentence i, as provided by GPTZero. We construct an empirical PDF that captures where AI probability mass is concentrated within the article.

Position Normalization. We normalize each sentence's position to the unit interval:

$$p_i = \frac{i-1}{n-1}, \quad i = 1, \ldots, n \tag{IA.9}$$

where $p_i \in [0,1]$ represents sentence i's relative position, with $p_1 = 0$ (first sentence) and $p_n = 1$ (last sentence).

PDF Weights. We define the empirical PDF weight for each sentence as its share of total AI probability:

$$w_i = \frac{a_i}{\sum_{j=1}^{n} a_j} \tag{IA.10}$$

By construction, $\sum_{i=1}^{n} w_i = 1$, so the weights form a proper probability distribution over positions. This formulation captures where AI-generated content is concentrated: if AI probability is high at the edges and low in the middle, the PDF will have mass concentrated near $p = 0$ and $p = 1$.

Position Regions. We partition the unit interval into three regions corresponding to article structure:

$$\text{Intro (Edge):} \quad p < 0.15 \tag{IA.11}$$

$$\text{Body:} \quad 0.15 \leq p \leq 0.85 \tag{IA.12}$$

$$\text{Conclusion (Edge):} \quad p > 0.85 \tag{IA.13}$$

The edge regions (intro and conclusion) comprise approximately 30% of positions, while the body comprises 70%.

C.3 Benchmark PDF Definitions

We define benchmark PDFs representing the idealized positional distribution for each usage category. Each benchmark assigns weights to positions, which are then normalized to sum to one.

Edge-Loaded AI Benchmark. AI probability mass is concentrated at the edges:

$$b_i^{\mathrm{EL}} \propto \begin{cases} 4 & \text{if } p_i < 0.15 \text{ or } p_i > 0.85 \\ 1 & \text{otherwise} \end{cases} \tag{IA.14}$$

The 4:1 weight ratio implies that edge positions receive four times the probability mass of body positions, yielding approximately 63% of total mass in the edge regions after normalization.

Body-Loaded AI Benchmark. AI probability mass is concentrated in the body:

$$b_i^{\mathrm{BL}} \propto \begin{cases} 3 & \text{if } 0.15 \le p_i \le 0.85 \\ 1 & \text{otherwise} \end{cases} \tag{IA.15}$$

The 3:1 weight ratio places approximately 87% of mass in the body region.

Uniform AI Benchmark. AI probability is uniform across all positions:

$$b_i^{\mathrm{U}} \propto 1 \quad \text{for all } i \tag{IA.16}$$

This yields a uniform distribution where each position receives equal weight.

Interspersed AI. This category cannot be characterized by a positional benchmark because its defining feature is frequent *transitions* between AI and human sentences rather than concentration at particular positions. We detect this pattern using the transition rate, defined below.

C.4 Distance Metric and Transition Rate

L1 Distance (Total Variation Distance). We measure the dissimilarity between an article's empirical PDF and each benchmark using the L1 distance, also known as total variation distance:

$$d(w, b) = \frac{1}{2} \sum_{i=1}^{n} |w_i - b_i| \tag{IA.17}$$

where $w = (w_1, \ldots, w_n)$ is the empirical PDF and $b = (b_1, \ldots, b_n)$ is a benchmark PDF (interpolated to match the article's sentence count). The factor of $\frac{1}{2}$ normalizes the distance to $[0, 1]$, where 0 indicates identical distributions and 1 indicates completely non-overlapping distributions.

For each article, we compute three distances:

- d_{EL}: Distance to Edge-Loaded AI benchmark
- d_{BL}: Distance to Body-Loaded AI benchmark
- d_{U}: Distance to Uniform AI benchmark

Transition Rate. To detect Interspersed AI, we compute the rate at which consecutive sentences switch between AI-likely and human-likely classifications:

$$\tau = \frac{1}{n-1} \sum_{i=2}^{n} \mathbf{1}\left[(a_i > 0.5) \neq (a_{i-1} > 0.5)\right] \tag{IA.18}$$

where $\mathbf{1}[\cdot]$ is the indicator function. A high transition rate indicates frequent alternation between AI and human sentences, consistent with iterative editing throughout the article.

C.5 Classification Algorithm

We classify articles using the following procedure:

1. **Filter:** We restrict to articles with AI coverage ≥ 15%, defined as the share of sentences with $a_i > 0.5$. Contents below this threshold have insufficient AI content for meaningful pattern classification.
2. **Check Interspersed AI:** If $\tau > 0.15$ and $d_{\text{U}} < 0.40$, the content is classified as *Interspersed AI*.
   - *Rationale:* High transition rate (> 15%) indicates frequent AI-human alternation. The constraint $d_{\text{U}} < 0.40$ ensures the article has sufficiently uniform positional distribution (ruling out cases where transitions occur but AI is still concentrated at edges or body).
3. **Assign by Minimum Distance:** If the content is not classified as *Interspersed AI*, then the content is classified based on the smallest distance:

$$\text{Category} = \arg\min_{c \in \{\text{EL,BL,U}\}} d_c \tag{IA.19}$$

   A content is assigned to *Edge-Loaded AI* if d_{EL} is smallest, *Body-Loaded AI* if d_{BL} is smallest, or *Uniform AI* if d_{U} is smallest.

Why Check Interspersed AI First? Interspersed AI produces a roughly uniform positional distribution similar to Uniform AI, since AI edits occur throughout the article. Without the transition rate check, these articles would be misclassified as Uniform AI. By checking for high transition rates first, we correctly identify the iterative editing pattern before defaulting to position-based classification.

C.6 Output Variables

The classification procedure yields the following variables for each article:

- `category`: Categorical variable ∈ {Uniform AI, Body-Loaded AI, Edge-Loaded AI, Interspersed AI}
- `dist_edge`: L1 distance to Edge-Loaded AI benchmark
- `dist_body`: L1 distance to Body-Loaded AI benchmark
- `dist_uniform`: L1 distance to Uniform AI benchmark
- `transition_rate`: Share of consecutive sentence pairs that switch AI/human classification
- `ai_coverage`: Share of sentences with AI probability > 0.5
- `mean_ai_prob`: Average AI probability across all sentences

The continuous distance measures allow for regression analysis examining how proximity to each benchmark relates to outcomes of interest, while the categorical variable enables comparison across discrete usage types.

C.7 Limitations

Several limitations merit acknowledgment. First, our benchmark PDFs reflect stylized patterns that may not capture all forms of human-AI collaboration. Second, the transition rate threshold for Interspersed AI (15%) and the distance threshold (0.40) are based on inspection of the data rather than theoretical derivation. Third, articles near decision boundaries may be sensitive to small changes in thresholds. Fourth, our method assumes that positional patterns are informative about authoring workflow, which may not hold for all article types or lengths.

Table IA1. AI-Generated Content and Sentiment Alignment: Robustness

This table replicates the sentiment alignment analysis from Table 5 using only content that received at least one comment. The dependent variable is *Comment Sentiment*$_{ij[t,t+10]}$, the mean sentiment of comments following content j about firm i over $[t, t+10]$. *AIProb*$_{ijt}$ is the estimated probability, ranging from 0 to 1, that content j covering firm i on day t is AI-generated. *Sentiment*$_{ijt}$ is the main text sentiment score. *Controls* includes the article characteristics, firm characteristics, and author history performance (standardized to zero mean and unit variance). Columns 1–2 report results for contents on Reddit's r/WallStreetBets; Columns 3–4 report results for articles published on Seeking Alpha. Standard errors are two-way clustered by stock and day; t-statistics are in parentheses. *p <.1; **p <.05; ***p <.01.

	SA		WSB	
	(1)	(2)	(3)	(4)
AIProb × Sentiment	-0.061	-0.099	0.147*	0.376***
	(-0.76)	(-1.19)	(1.74)	(2.61)
AIProb	-0.001	0.001	-0.008**	-0.022***
	(-0.21)	(0.39)	(-2.24)	(-2.75)
Sentiment	0.265***	0.269***	0.148***	0.105***
	(11.26)	(10.28)	(8.69)	(3.88)
Length	0.003***	0.002***	-0.002***	-0.001
	(9.17)	(5.09)	(-4.15)	(-1.04)
Complexity	0.000	0.001	0.000	0.000
	(1.16)	(1.60)	(-0.58)	(0.07)
Fog Index	0.000	0.002***	-0.001	0.000
	(0.91)	(3.50)	(-1.14)	(0.04)
Graphical	0.001***	0.001***	0.000	0.000
	(4.09)	(3.45)	(0.68)	(-0.21)
Quantitative	0.001***	0.001**	0.001**	0.002
	(3.16)	(1.97)	(2.43)	(1.64)
Size	0.003	0.005	-0.002	-0.001
	(0.96)	(1.39)	(-0.37)	(-0.08)
BM	-0.001	-0.001	-0.002	-0.003
	(-1.00)	(-0.97)	(-1.65)	(-0.91)
IO	0.001	0.001	0.000	-0.002
	(0.75)	(0.47)	(-0.04)	(-0.48)
News	-0.001*	-0.001	-0.001	-0.002
	(-1.71)	(-1.39)	(-0.71)	(-1.19)
Analysts	0.000	0.000	0.001	0.002
	(0.00)	(-0.35)	(0.53)	(0.84)
Author Article	0.000	0.000	0.000	-0.001
	(0.24)	(-0.05)	(-0.60)	(-0.66)
Author AI Article	0.000	0.000	0.000	-0.001
	(-0.32)	(-0.92)	(-0.97)	(-1.01)
Author Comments	0.000	0.000	0.001*	0.001
	(0.11)	(0.42)	(1.93)	(0.84)
Submissions			0.012***	0.011***
			(12.26)	(3.09)
Day FE	X	X	X	X
Stock FE	X	X	X	X
Author FE		X		X
Obs.	43,937	43,937	37,487	37,487
Adj. R^2 (%)	14.0	19.1	5.2	51.4

Table IA2. WallStreetBets Analysis: Alternative Samples

This table presents replication results of Tables 3–6 using alternative WallStreetBets samples. Panel A includes all Messages (submissions and comments with fewer than 50 words) and Panel B excludes messages that were later removed or deleted. Columns 1-3 examine article-level AI adoption; Columns 4-5 examine the association between AI content and user responses; Columns 5–6 relate *AIProb* to cumulative abnormal returns. Standard errors are two-way clustered by stock and day; *t*-statistics are shown in parentheses. $^{*}p$ <.1; $^{**}p$ <.05; $^{***}p$ <.01.

A. Including Short Messages (Fewer Than 50 Words)

	AIProb			Cmmt. Sentiment	CAR[1,5]	CAR[6,10]
	(1)	(2)	(3)	(4)	(5)	(6)
Unfamiliarity	0.017					
	(1.53)					
News		-0.005		0.000	-0.171	-0.084
		(-0.61)		(-1.58)	(-0.70)	(-0.38)
Analysts		0.032***		0.000	-0.158	0.031
		(3.33)		(-1.32)	(-0.34)	(0.10)
IntenseRetailBuy			0.105*			
			(1.85)			
AIProb x Sentiment				0.137***	-0.686	-0.333
				(2.87)	(-0.90)	(-0.38)
Sentiment				0.018***	0.021	-0.005
				(25.02)	(1.51)	(-0.44)
AIProb				-0.009***	-0.785**	0.496
				(-4.26)	(-2.45)	(1.53)
Word Count > 50	2.675***	2.677***	2.572***	-0.002**	-0.097	-0.116
	(28.10)	(27.99)	(27.49)	(-7.266)	(-1.22)	(-1.13)
Day FE	X	X	X	X	X	X
Stock FE	X	X	X	X	X	X
Author FE	X	X	X	X	X	X
Obs.	949,674	949,674	901,387	949,674	942,514	942,514
Adj. R^2 (%)	46.6	46.6	46.0	10.3	28.9	30.9

B. Excluding Deleted Messages

	AIProb			Cmmt. Sentiment	CAR[1,5]	CAR[6,10]
	(1)	(2)	(3)	(4)	(5)	(6)
Unfamiliarity	0.015					
	(0.10)					
News		-0.219		0.000	-0.313	-0.171
		(-1.44)		(-0.47)	(-1.09)	(-0.66)
Analysts		0.094		0.002*	0.060	0.352
		(0.43)		(1.66)	(0.10)	(0.84)
IntenseRetailBuy			1.686**			
			(2.33)			
AIProb x Sentiment				0.163***	-0.146	-0.472
				(2.62)	(-0.33)	(-1.22)
Sentiment				0.065***	-0.010	0.000
				(6.50)	(-0.13)	(-0.01)
AIProb				-0.010**	-1.216*	-0.111
				(-2.98)	(-1.73)	(-0.17)
Day FE	X	X	X	X	X	X
Stock FE	X	X	X	X	X	X
Author FE	X	X	X	X	X	X
Obs.	59,104	59,104	56,263	59,104	58,565	58,565
Adj. R^2 (%)	51.7	51.7	53.2	0.2	26.2	28.0