Sycophancy or Information Extraction? Evidence from Real-world Human–GAI Interactions

July 14, 2026

Abstract

This study examines sentiment adaptation between human users and generative artificial intelligence (GAI) in financial markets, using proprietary data from a leading GAI platform in China. We propose that GAI's adaptation reflects strategic extraction and integration of substitutive human information. Because of updating frictions, human adaptation is more limited than GAI adaptation. Consistent with this view, we first document pronounced asymmetry: GAI sentiment responds strongly to user inputs, whereas user sentiment responds only modestly to GAI. GAI adaptation is also selective. It is stronger when user sentiment is more informative, when the model's domain knowledge is more limited, when user and GAI information is more independent, and when it can better interpret user inputs, e.g., with search-based retrieval. These dynamics have meaningful economic consequences: greater adaptation is associated with more accurate beliefs, stronger informational advantages for GAI, and improved trading performance. Overall, our evidence suggests that GAI functions as an interactive intermediary that extracts user information to generate more informative responses while simultaneously shaping user beliefs, which seems to enhance, rather than undermine, market efficiency.

Keywords: Generative Artificial Intelligence; Sycophancy; Information Aggregation; Investor Sentiment; Belief Update

JEL Codes: D83; G14; G41; O33

1 Introduction

Recent advances in generative artificial intelligence (GAI) are transforming how investors acquire, interpret, and act on information (e.g., Ecker et al. (2026)). Unlike traditional information technologies, large language models (LLMs) operate through interactive dialogue: users iteratively exchange information with the model, and both parties adjust their responses throughout the conversation. As a result, belief formation increasingly occurs within private human–GAI interactions rather than solely through public information environments. Understanding the dynamics of these interactions is therefore critical for evaluating the informational and market consequences of GAI.

A growing literature argues that LLMs exhibit sycophantic behavior by aligning their responses with users' stated beliefs or preferences, even when those beliefs are incorrect (e.g. Perez et al. (2023); Sharma et al. (2024); Chen et al. (2025); Balasubramaniam and Sabat (2026)). This evidence raises concerns that GAI may amplify biases and distort belief formation. However, the observed GAI's adapation to users needs not arise solely from mechanical agreement-seeking behavior. *We explore the possibility that, through conversation adaptation, GAI can actually improve information accuracy in financial markets, where the pursuit of financial returns is high order.*[1]

We develop a simple framework in which human users and GAI receive noisy signals about the underlying fundamental firm value, but their signals differ in precision. Specifically, in many financial settings, even when users and GAI systems evaluate the same underlying fundamental, they possess different but partially overlapping informational advantages. Users often have contextual knowledge that is unavailable to the model, whereas GAI systems have comparative advantages in processing large-scale public information and historical patterns. Such partial overlap provides a basis for mutual belief updating. In such environments, conversational adaptation may emerge because GAI strategically extracts and integrates human information during interaction. Human users are subject to frictions that limit their updating from GAI-generated information

[1]Throughout this paper, we use the term "adaptation" to refer to the extent to which one party's stance moves toward the other's during human–GAI interactions.

(e.g., bounded rationality, limited attention, cognitive processing costs etc.), whereas GAI systems are explicitly designed to condition on user inputs. In equilibrium, GAI may optimally adapt towards users not because agreement itself is intrinsically valuable, but because adaptation facilitates the extraction of human information.

This framework generates three predictions. First, adaptation should be asymmetric, with GAI responding more strongly to users than users responding to GAI. Second, GAI adaptation should be selective. It should increase when GAI's relative informational advantage is smaller, that is, when user information is more informative and when the model's internal knowledge is weaker. It should also increase when users' and GAI's information is more independent, leaving more room for GAI to learn from user inputs, and when retrieval-based search helps GAI better distinguish signal from noise in those inputs. Third, if the asymmetric adaptation reflects information extraction rather than pure sycophancy, stronger adaptation should result in higher information precision for both users and GAI.

To test these predictions, we leverage a proprietary dataset of large-scale human–GAI real-world interactions from a leading Chinese GAI platform and focus on firm-specific financial discussions. Ecker et al. (2026) use similar data to evaluate investors' use of AI. Our final sample comprises roughly 50,000 conversation rounds across more than 14,000 multi-round dialogues linked to over 2,000 publicly listed firms over the period of September to December 2024, allowing us to track conversational dynamics at a granular level. Using a dynamic adjustment framework estimated via system GMM (Wintoki et al., 2012) and sentiment to reflect belief (e.g., Baker and Wurgler (2007); Tetlock (2007)), we document pronounced asymmetry in adaptation. GAI sentiment responds strongly to user sentiment: the estimated adjustment coefficient is close to 0.87, indicating that a one-unit increase in user sentiment is associated with a 0.87-unit increase in GAI tone. By contrast, user sentiment responds only modestly to prior GAI sentiment, with an adjustment coefficient of approximately 0.10. This asymmetry is robust across specifications and subsamples.

We next examine whether this adaptation reflects mechanical sycophancy which does

not depend on the state or selective incorporation of user information. If adaptation were driven by sycophancy, it should arise unconditionally across states. Instead, we find that adaptation is highly state dependent. First, GAI exhibits significantly stronger adaptation when user sentiment is more informative, measured in two ways: (1) whether the direction of user tone aligns with subsequent firm performance, and (2) whether users are financially sophisticated based on their use of professional financial terminology. Second, GAI adaptation varies with the precision of its own domain knowledge. We gauge the precision using the coverage by media available during the model's training period, because GAI are commonly trained on large-scale web corpora that include news content (Brown et al., 2020). More coverage implies a higher precision. When there is less coverage about a firm, we find that GAI adapts more strongly to user inputs, consistent with a greater reliance on external information when GAI's internal knowledge is limited. Third, adaptation is stronger when GAI and users exhibit greater disagreement in their initial sentiment, consistent with larger updating when the two signals are more independent. Fourth, GAI exhibits stronger adaptation in interactions where search is activated, consistent with retrieval based search helping GAI evaluate the precision of user signals. Taken together, these findings are more consistent with selective information extraction than with purely mechanical sycophancy.

Next, we link adaptation to economic outcomes. When GAI adaptation increases by 10% (relative to our baseline estimation, 0.87), users are 5.8% more likely to return to the GAI platform to initiate a new conversation, 7.2% less likely to update sentiment in the wrong direction, and the final-round GAI response is 5.3 percentage points more likely to be more accurate than the human user's final-round message. Moreover, GAI's adaptation is positively associated with subsequent investment returns evaluated based on either users' or GAI's sentiment in the final round. A 10% increase in GAI alignment relative to the baseline value of 0.868 corresponds to a 1.5% increase in trading profits evaluated based on users' final sentiment over the next 20 trading days and a 1.0% increase in trading profits based on the GAI's final sentiment over the same window. Taken together, these results support our prediction that GAI adaptation reflects an

endogenous information extraction process and doing so improves information precision resulting in higher trading profits.

Our study contributes to the emerging literature on GAI in capital markets in several ways. First, we add to the growing debate on AI sycophancy. Prior research in computer science and economics argues that large language models often accommodate user beliefs or preferences, raising concerns that GAI may amplify biases rather than provide independent information (e.g., Perez et al. (2023); Sharma et al. (2024); Chen et al. (2025); Balasubramaniam and Sabat (2026)). In contrast, our evidence suggests that observed adaptation in the real-world financial market is selective and state-dependent rather than purely mechanical. Specifically, GAI adapts more strongly when user inputs appear more informative, when the model's own information environment is weaker, and when retrieval-based search improves the interpretation of user inputs. These findings suggest that a substantial component of observed adaptation reflects strategic extraction and integration of human information rather than indiscriminate agreement.

Second, we conceptualize GAI not as a static information provider, but as an adaptive intermediary whose responses evolve endogenously through interaction with users. While prior studies primarily examine the informational content or investment implications of standalone GAI outputs (e.g., Ecker et al. (2026); Blankespoor et al. (2026b)), we focus on the dynamic process through which humans and GAI jointly form beliefs during multi-round conversations. In our framework, both parties possess information about underlying firm fundamentals, and conversational adaptation emerges through iterative interaction and learning. This perspective shifts attention from static AI recommendations to endogenous human–GAI information integration.

Third, our study contributes more broadly to the literature on information intermediation and market efficiency. A central insight of information economics is that economically relevant information is often dispersed across market participants and costly to aggregate (e.g., Hayek (1945); Grossman and Stiglitz (1980)). Traditional information intermediaries in capital markets such as analysts, financial media, auditors, and rating agencies facilitate price discovery by collecting, processing, and disseminating in-

formation to investors (e.g., Healy and Palepu (2001); Beyer et al. (2010); Bushee et al. (2018)). In contrast, conversational GAI operates through iterative interaction in which the system simultaneously learns from users while shaping users' beliefs. Our findings suggest that GAI may therefore function as an adaptive information intermediary that facilitates the incorporation of dispersed and decentralized human information into market beliefs and investment decisions. This mechanism differs from conventional forms of information production because information aggregation emerges endogenously through interaction between human and machine intelligence rather than through static disclosure or unilateral analysis. Consistent with this view, we show that stronger conversational adaptation is associated with more accurate beliefs and more profitable trading signals. More broadly, our evidence suggests that the economic effects of GAI depend not only on model quality, but also on how effectively conversational systems extract, interpret, and integrate value-relevant information from market participants.

2 Background and Theoretical Framework

2.1 Human–GAI Adaptation in Financial Decision-Making

In financial settings, investors increasingly rely on GAI not only for information acquisition but also for information integration (Ecker et al., 2026; Blankespoor et al., 2026a). This shift relocates a substantial portion of belief formation from public information environments to private, GAI-mediated interactions, where users iteratively refine their views through dialogue with the model.

A growing body of work highlights a feature of such interactions: GAI tends to align its responses with users' stated views, a behavior commonly referred to as sycophancy (Cheng et al., 2025b, 2026; Chandra et al., 2026). Recent public reactions suggest that users are increasingly aware of GAI's tendency to align with human interlocutors. In April 2025, OpenAI rolled back a GPT-4o update after acknowledging that the model had become overly flattering and agreeable, a behavior the company explicitly described as "sycophantic." OpenAI attributed this issue partly to excessive reliance on short-

term user feedback, which led the model to generate overly supportive but disingenuous responses.[2] These concerns motivate us to explore how conversational adaptation affects belief formation and decision-making in financial settings.

On one hand, GAI adaptation can reinforce users' prior beliefs and contribute to biased or distorted inference. Consistent with this perspective, experimental evidence shows that exposure to GAI responses that are consistent with users' own views increases users' confidence in those views, reduces their willingness to revise judgments, and may influence subsequent behavior (Cheng et al., 2026). More broadly, conversational GAI can exert persuasive effects through interactive dialogue, with the potential to shape beliefs (Naddaf, 2025). Formal models further suggest that repeated interaction with a sycophantic GAI can generate feedback loops in belief updating, amplifying initial bias even for otherwise rational users (Chandra et al., 2026).

On the other hand, the literature also points to an important conceptual ambiguity, particularly in financial settings characterized by dispersed and heterogeneous information. Prior studies show that AI systems can generate highly accurate outputs when underlying information is transparent and abundant (e.g., Cao et al. (2024)). Similarly, algorithmic—or AI-assisted—analysts produce a more balanced distribution of buy, hold, and sell recommendations than their human counterparts (Coleman et al., 2022). In the GAI era, the capabilities of AI models have expanded substantially, enabling their use across a wide range of financial tasks (e.g., Blankespoor et al. (2026a); Cao et al. (2024); Cheng et al. (2025a); Bertomeu et al. (2025); Sheng et al. (2026); Blankespoor et al. (2026b); Chang et al. (2026); Bradshaw et al. (2026)). These findings suggest that GAI systems are capable of effectively processing rich information environments, raising the possibility that adaptation to user inputs may, in some cases, reflect efficient information extraction rather than bias. In many real-world settings, particularly in investment contexts, human users are not merely expressing preferences or emotions (e.g., "I like A better than B"). Instead, they may possess and convey contextual, tacit, or early-stage information that is not directly observable to the model—for example, qualitative insights

[2] https://openai.com/index/sycophancy-in-gpt-4o/

about firm-specific developments, local market conditions, or the relevance of particular signals. Moreover, because financial decisions involve direct wealth consequences, users often have strong incentives to process information carefully and act cautiously. As a result, user inputs need not be purely noisy or biased; they may contain information. From this perspective, GAI'S adaptation to user sentiment need not reflect the reinforcement of prior biases. Rather, it may arise from the extraction and integration of relevant information from users, particularly information that is not otherwise available to the model.

2.2 Hypothesis Development

2.2.1 Asymmetric Adaptation

Consistent with existing studies (e.g., Bertomeu et al. (2025); Cheng et al. (2025a); Bradshaw et al. (2026); Li et al. (2024)), we view GAI and humans possess distinct information about the same underlying fundamentals: GAI systems are well suited to synthesizing large-scale public information, retrieving historical analogs, and structuring unstructured data, whereas Human users may possess contextual, tacit, or early-stage information that is not directly observable to GAI. We therefore conceptualize GAI adaptation as the result of belief updating: GAI moves toward users when it extracts information from their inputs.

In our framework, the initial sentiment expressed by each party conveys a private, noisy signal of the firm's underlying fundamental value, with human and GAI signals differing in precision. Observing a counterpart's sentiment reveals information about their private signal, thereby inducing a belief update. This dynamic renders the interaction value-enhancing: multi-round dialogue serves as a mechanism to aggregate dispersed information, reducing the variance of the fundamental estimate. Crucially, sentiment adaptation between the user and the GAI is not mechanically imposed; rather, it emerges endogenously. As both agents iteratively update their beliefs using signals that are independent conditional on the fundamental, their posterior estimates naturally converge.

Although both parties can, in principle, learn from each other, their updating behavior is not symmetric. Human updating might be shaped by frictions, including confirmation bias, anchoring, limited attention, and lack of trust in algorithmic outputs (e.g., Rabin and Schrag (1999); Barberis and Thaler (2003); Jussupow et al. (2024); De Freitas et al. (2023); Leib et al. (2024)). As a result, humans might revise their beliefs gradually and incompletely. GAI systems, in contrast, are explicitly designed to extract user information and incorporate it into their responses, making them more responsive to external inputs. Therefore, in the presence of conditionally independent information, GAI adapts more strongly to human sentiment than humans adapt to GAI. We argue that this asymmetry reflects differences in updating mechanisms rather than purely behavioral "agreement-seeking" motives. This line of argument leads to our first hypothesis:

Hypothesis 1. ***(H1).*** *GAI sentiment responds more strongly to prior human sentiment than human sentiment responds to prior GAI sentiment.*

2.2.2 Relative Information Precision

In our framework, the magnitude of the GAI's belief update is governed by the relative precision of the human user's signal. When the user's signal is highly precise compared to the GAI's own information, the GAI undergoes a substantial belief update. Conversely, if the user's signal is noisy—or the GAI's prior information is highly precise—the update is muted, directly reflecting the lower informational value of the external input. Based on this line of reasoning, we formulate our second hypothesis:

Hypothesis 2a. ***(H2a).*** *GAI adaptation toward user sentiment is stronger when user-provided information is more precise.*

Hypothesis 2b. ***(H2b).*** *GAI adaptation toward user sentiment is stronger when the model's internal knowledge is less precise.*

2.2.3 Information Overlap

A further cross-sectional prediction is that GAI adaptation is stronger when there is greater initial disagreement between the GAI and the user, as this divergence implies

greater signal independence. In our setting, both parties hold priors regarding the same underlying firm fundamental but rely on distinct signals. When the GAI and the user begin with similar sentiments, the user's input likely overlaps with information already embedded in the model's prior, leaving less room for updating. Conversely, substantial initial disagreement suggests that the user's signal contains novel information relative to the model's own. In the limiting case where the two signals are fully independent, the interaction operates as a standard Bayesian updating process where each party learns from the other. Therefore, provided the model treats user sentiment as informative rather than mechanically echoing it, greater initial disagreement should induce a stronger belief update toward the user. This reasoning leads to our third hypothesis:

Hypothesis 3. ***(H3).*** *GAI adaptation toward user sentiment is stronger when the two signals are initially more independent.*

A distinctive feature of modern GAI systems is their ability to augment interactions with retrieval-based search and other external information tools. Importantly, these tools not only provide additional information, they can also improve the model's ability to evaluate the information content embedded within a user's prompt. Institutional context supports this mechanism: according to our discussions with system engineers at our partner GAI platform, the model's decision to activate an external search is largely determined by the precision of the user's input.[3] In short, by enhancing the system's ability to evaluate user informativeness, retrieval mechanisms enable GAI to extract and synthesize user-provided insights more effectively This leads to our fourth hypothesis:

Hypothesis 4. ***(H4).*** *GAI's adaptation towards user sentiment is more pronounced when information-search tools are activated.*

[3]For instance, consider a user stating that *"NVIDIA's recent supplier commentary suggests stronger demand than analysts currently project."* Retrieval does not inherently improve the precision of the user's signal itself. Instead, it empowers the GAI system to cross-reference the claim by querying news articles, extracting earnings call transcripts, and auditing recent disclosures. By synthesizing these external sources with the user's input, the model can systematically evaluate whether the claim is plausible, determine if empirical evidence supports it, and assess whether the user possesses incremental information not yet reflected in public consensus.

2.2.4 Economic Outcomes

Through iterative interaction, both the GAI and the human user update their beliefs, leading to more accurate posterior estimates of the underlying fundamental. Consequently, executing trades based on these refined posteriors yields greater expected profitability. Our fifth hypothesis is thus stated as follows:

Hypothesis 5a. ***(H5a).*** *Greater adaptation by the GAI is associated with its more accurate posterior estimates of the firm's fundamental.*

Hypothesis 5b. ***(H5b).*** *Greater adaptation by the GAI is associated with higher trading profits based on its posterior estimates.*

3 Data and sample selection

3.1 Sample Construction

To test our hypotheses, we need to directly observe the actual conversations between users and the GAI, as well as the subsequent consequences. Thus, we collaborate with a leading Chinese generative GAI platform and rely on the real-world anonymized human–GAI interaction data to conduct empirical analysis.[4] We collected a sample that contains all real-world users' conversations at the platform spanning September to December 2024. For each interaction, we observe the complete sequence of question–answer exchanges, including (i) timestamps and anonymized user identifiers, (ii) the full prompt and response texts, (iii) conversation identifiers that allow us to reconstruct multi-round interactions, and (iv) system-level metadata on the underlying model, tools, and status flags. In our empirical analyses, we define a conversation as a single, uninterrupted interactive session between the user and the GAI. Each conversation comprises multiple rounds of interaction, with a single round defined as a user prompt followed by the corresponding

[4] Although lab experiments based on fake users may also offer similar observations on GAI responses, the prompts and tasks may not be representative of real-world usage, thus compromising external validity. Moreover, even if researchers later give GAI the same information investors had at the time, the resulting prediction is generated only ex-post, and the market responses may have already occurred. It therefore cannot show that investors could actually use that prediction to earn money, because there could be look-ahead bias.

GAI response

We focus on capital market setting, where information is rich and readily available. It also allows us to evaluate the economic consequences of GAI–human interactions by benchmarking against the market outcomes. At the same time, it is a professional context with direct wealth implications, in which GAI's sycophantic behavior may be particularly consequential. We specifically examine investment-related GAI-human interactions—namely, discussions of publicly listed firms (Ecker et al., 2026; Blankespoor et al., 2026a). To do so, we match GAI–human conversations to publicly listed firms by identifying firm references (i.e., firm names and stock tickers) within the conversation texts. This step yields roughly 250,000 conversations and over one million conversation rounds, covering nearly 3,800 publicly listed firms that represent more than 70% of China's total stock market capitalization.[5]

We then apply a sequence of data-filtering criteria. First, we drop conversations containing duplicate firm identifiers to ensure a strict one-to-one mapping between dialogues and firms, thereby sharpening the precision of our firm-level inferences. Within this step, we also remove any conversation rounds flagged with generation errors (e.g., `status = "error"`). Second, we exclude multimodal interactions involving non-text inputs (i.e., images), observations with missing or empty text fields, and outliers with extremely short or excessively long prompts (specifically, those outside a range of 5 to 1,000 characters). These restrictions ensure that the final sample consists of textual exchanges that can be reliably measured and analyzed, and that are highly likely to reflect natural, informative interactions. Third, we exclude conversations centered on purely mechanical text-processing tasks such as translation, polishing, rewriting, and content expansion, since these interactions do not capture genuine investment analysis. This process yields a final sample of approximately 100,000 conversation rounds spanning roughly 32,000 distinct conversations across 2,700 firms.

Unlike Ecker et al. (2026) and Blankespoor et al. (2026a), we focus on GAI–human exchanges, which requires sufficiently long conversational sequences. Hence, we retain

[5]Our sample differs from that in Ecker et al. (2026), likely because the two studies rely on data from different GAI platforms.

only multi-round dialogues. We further employ a fine-tuned Chinese BERT classifier to screen out dialogues unrelated to capital markets.[6] These deleted conversations include, for example, customer support requests (e.g., account setup or troubleshooting), product usage or after-sales issues, and general corporate inquiries such as recruitment, office locations, or internal operations. Our final sample consists of roughly 50,000 conversation rounds from more than 14,000 multi-round conversations, covering about 2,300 firms. For the GMM analyses, which require sufficient lagged variables, the sample is further reduced to approximately 8,000 conversation rounds from roughly 4,000 conversations, covering about 1,200 firms. Table 1 describes in detail our sample selection process. Appendix 9 presents anonymized examples that illustrate investment-related conversations.

3.2 Summary Statistics

We compare the firms in our final human–GAI conversation sample with the universe of listed firms in the Chinese stock market at the end of 2024 in Table 2. Compared to the average Chinese listed firm, the firms in our sample are larger, older, more leveraged, and attract greater analyst coverage. However, their financial and stock performance metrics do not differ significantly. The book-to-market ratio is marginally higher for our sample firms, with the difference statistically significant at the 10% level. Overall, our sample captures firms characterized by greater market visibility and richer public information environments, consistent with the types of equities that typically attract retail investor attention. In subsequent analyses, we control for these variations across specifications.

4 Research Design

Each conversation in our sample, c, consists of multiple rounds. Each conversation round r begins with a user prompt followed by the GAI's response. Accordingly, the GAI's sentiment in each round can be modeled as a weighted combination of its prior stance

[6] To identify conversations related to capital-market information, we train a supervised text classifier using a manually labeled sample of user prompts randomly drawn from the original conversations. We split the labeled data into training, validation, and test sets and fine-tune a BERT-based classifier, selecting the final specification based on validation performance. The resulting model performs well out of sample, with accuracy and macro F1 both around 90%.

and the new information provided by the user. Following (Huang et al., 2023), we gauge sentiment using tone, calculated based on finance-domain FinBERT model. Specifically, we use FinBERT to estimate the probabilities that each sentence has a positive, negative, or neutral tone, where tone is defined as the positive probability minus the negative probability.

Before estimating the dynamic adjustment models, we demean both GAI and human user tone at the firm and date levels. Specifically, we remove persistent firm-specific sentiment and date-specific common shocks from each tone measure. We denote the resulting demeaned variables by ($\widetilde{Tone}$), and all tone variables in the equations below refer to these demeaned values. Combined with the adjacent-round structure of the model, this transformation allows the estimated coefficients to capture sequential tone adjustments within conversations, net of systematic firm- and date-level variation. The following equations formalize the information aggregation process within our theoretical framework. Specifically, the GAI system updates its belief by synthesizing its prior knowledge with the new information extracted from the user's input, weighting both signals. Empirically, we rely on text-derived sentiment to capture and quantify this information content (Loughran and McDonald, 2011; Chang et al., 2026).

$$\widetilde{Tone}_{c,r}^{GAI} - \widetilde{Tone}_{c,r-1}^{GAI} = \lambda_{GAI}\left(\widetilde{Tone}_{c,r}^{Human} - \widetilde{Tone}_{c,r-1}^{GAI}\right) + \varepsilon_{c,r}^{GAI}, \tag{1}$$

which can also be written as

$$\widetilde{Tone}_{c,r}^{GAI} = (1-\lambda_{GAI})\widetilde{Tone}_{c,r-1}^{GAI} + \lambda_{GAI}\widetilde{Tone}_{c,r}^{Human} + \varepsilon_{c,r}^{GAI}. \tag{2}$$

Equation (1) closely mirrors the partial adjustment models in the dynamic capital structure literature (Flannery and Rangan, 2006). A critical distinction, however, lies in the nature of the target: while traditional corporate finance models assume a well-defined, steady-state target toward which firms converge over time, the 'target' in our setting is ex ante unobservable and endogenous, evolving dynamically with the user's inputs across successive conversation rounds. In the empirical analyses that follow, we operationalize

the above two equations interchangeably.

Similarly, human users weigh GAI's response from the preceding round and their own prior in the following way:

$$\widetilde{Tone}_{c,r}^{Human} - \widetilde{Tone}_{c,r-1}^{Human} = \lambda_{Human}\left(\widetilde{Tone}_{c,r-1}^{GAI} - \widetilde{Tone}_{c,r-1}^{Human}\right) + \varepsilon_{c,r}^{Human}, \tag{3}$$

which is equivalent to:

$$\widetilde{Tone}_{c,r}^{Human} = (1-\lambda_{Human})\widetilde{Tone}_{c,r-1}^{Human} + \lambda_{Human}\widetilde{Tone}_{c,r-1}^{GAI} + \varepsilon_{c,r}^{Human}. \tag{4}$$

Within this framework, λ_{GAI} captures the extent to which GAI adapts to the human user, while λ_{Human} captures the degree to which the human user adapts to the GAI. A larger adaptation coefficient implies greater convergence toward the counterpart, whereas a smaller coefficient implies greater persistence in the agent's own prior belief. The terms ε_{cr}^{GAI} and ε_{cr}^{Human} denote idiosyncratic shocks to the GAI tone and human user tones, respectively.

Empirically estimating the dynamics of tone adaptation is challenging because GAI–human conversations exhibit serial dependence: prior exchanges shape subsequent responses, and earlier dialogue history continues to influence current interactions. To address this issue, we estimate the above equations using a dynamic GMM framework. This approach is well-suited to our setting because it accommodates panel data with serial correlation and captures the iterative nature of dialogue, in which both GAI and human users continuously update their prior stance and respond to the counterpart's feedback from the previous round. We therefore model within-conversation adaptation using the equations above as reduced-form representations.[7]

[7]While Equations (1) and (3) could be estimated using pooled OLS or fixed-effects (FE) estimators as well, such approaches may yield biased and inconsistent estimates. The primary econometric challenge is *dynamic endogeneity* (Wintoki et al., 2012), which arises when current explanatory variables are shaped by past realizations of the dependent variable. In our setting, current GAI and human tones may depend on prior-round tones, while the counterpart's current tone may also respond to the same conversational history. This dynamic structure makes standard fixed-effects estimators biased (Nickell, 1981). We therefore estimate dynamic panel GMM specifications based on the difference-GMM and system-GMM approaches of Arellano and Bond (1991) and Blundell and Bond (1998), following their application to corporate finance settings in Wintoki et al. (2012). Nevertheless, we provide robustness of our main

Following Wintoki et al. (2012), we estimate both *difference GMM estimator* and *system GMM estimator*. The former eliminates conversation-specific, time-invariant effects by first-differencing the model and using lagged levels of the endogenous variables as instruments for the equations in differences.[8] However, the effectiveness of difference GMM estimator hinges on the strength of the correlation between the instruments (lagged levels) and the endogenous regressors (differences). In our context, both GAI and human tones exhibit substantial persistence across conversation rounds. When variables are highly persistent, lagged levels become weak instruments for the variables in differences, potentially leading to finite-sample bias and a loss of precision. Therefore, we rely on the latter, the *system GMM estimator*, as our primary specification, but still present the results of *difference GMM estimator* alongside. System GMM addresses the weak-instrument problem by estimating a system of equations in both first differences and levels. In this framework, lagged levels are used as instruments for the differenced equations, and lagged differences are used as instruments for the level equations.

To ensure the validity of our GMM instruments, we perform two standard diagnostic tests. First, we use the Hansen J-test of overidentifying restrictions to evaluate the overall validity of the instrument set (Hansen, 1982). Second, we conduct the Arellano–Bond test to evaluate serial correlation in the differenced residuals (Arellano and Bond, 1991). While first-order serial correlation ($AR(1)$), is expected in the differenced residuals by construction, the absence of second-order serial correlation ($AR(2)$) supports the assumption that the error terms in the original level equations are not serially correlated and that the lagged instruments are valid. The system GMM estimator further follows Blundell and Bond (1998), who augment the differenced equations with level equations instrumented by lagged differences.

results in Table A2 and our inference holds.

[8] In difference GMM, the "equations in differences" refer to the first-difference of the original level equations, which remove time-invariant conversation-specific effects. Because the differenced lagged dependent variable remains correlated with the differenced error term, sufficiently lagged levels of the endogenous variables are normally used as internal instruments for the endogenous regressors.

5 Empirical Results

This section presents our empirical results. Section 5.1 visually demonstrates the within-conversation sentiment dynamics. Section 5.2 formally tests H1, i.e., whether GAI adapts more strongly to human users than vice versa in conversations. Section 5.3 investigates the underlying mechanisms–our H2–H4. Finally, Section 5.4 evaluates the economic outcomes of GAI adaptation–our H5.

5.1 Descriptive Evidence

Figure 1 Panel A plots the average tone across successive conversation rounds, bounded by 95% confidence intervals. Because fewer conversations persist into the later stages, the number of observations drops sharply in subsequent rounds. This sample attrition explains the widening of the confidence intervals as a conversation progresses. As shown, GAI responses are, on average, more positive than user messages, while user tone is roughly neutral. We also observe a downward trend in the average tone of GAI over the course of a conversation, while user tone exhibits a slight decline. Figure 1 Panel B plots the median tone in each round using dashed lines, which exhibit a pattern similar to the mean series. The tendency for GAI tone to move toward user tone is even more pronounced in the median series, whereas user tone remains relatively stable as a conversation progress.

Table 4 confirms this finding with a univariate test. GAI tone decreases from the initial round to the final round: its mean declines from 0.381 to 0.360, with a difference being statistically significant at the 1% level. The median declines from 0.168 to 0.048, with the difference also being statistically significant at the 1% level. By contrast, user tone remains much closer to neutral. Its mean increases slightly from 0.023 to 0.032, with a mean difference of 0.008 that is statistically significant at the 1% level. However, its median is virtually unchanged, increasing only from 0.000607 to 0.000632; the corresponding median difference is 2.50×10^{-5} and is not statistically significant. These results provide preliminary evidence that the tone convergence is primarily driven by the decline

in GAI tone, while the central tendency of user tone changes little over the course of the interaction. These findings provide preliminary evidence consistent with Hypothesis 1 (H1), supporting the predicted asymmetric adaptation between human users and the GAI system

5.2 GAI-user Tone adaptation

To formally test H1, Table 5 reports the GMM estimation results. Columns (1) and (2) present the estimates of GAI'S adaptation to users (i.e., Equation (2)), while columns (3) and (4) report the corresponding analysis of user's adaptation to GAI (i.e., Equation (4)). We report results from both Difference GMM and System GMM estimators. We assess the reliability of our GMM estimation by reporting two standard diagnostic checks. The Hansen J-test examines whether the set of instruments as a whole is valid (Hansen, 1982), while the Arellano–Bond test checks for serial correlation in the first-differenced residuals (Arellano and Bond, 1991). As expected, first-order correlation ($AR(1)$) appears in the differenced residuals; importantly, the absence of second-order correlation ($AR(2)$) suggests that the original-level error terms are not serially correlated, justifying the use of lagged instruments. Our system GMM specification follows Blundell and Bond (1998), combining differenced equations with level equations instrumented by lagged differences to improve efficiency and reduce finite-sample bias. As shown, these diagnostic statistics support the validity of our tests.

Note that the estimated adjustment magnitude is substantially larger for GAI responding to human sentiment in the current round than vice versa. The economic magnitude is also meaningful. In the System GMM specification (Column (2)), the coefficient of 0.868 indicates that GAI places approximately 86.8% weight on user sentiment and only 13.2% on its own prior stance (see Equation (2)). In contrast, the adjustment speed on the human side is much smaller. The corresponding coefficient of 0.105 implies that users place only 10.5% weight on GAI responses, while their priors account for 89.5% of the weight (see Equation (4)).

The magnitude remains economically meaningful even after accounting for the dis-

persion of tone in GAI and users. Based on the standard deviations reported in Table 3, a one-standard-deviation increase in contemporaneous user tone (0.307) leads to a 0.266 increase in GAI tone within the same round (0.868 × 0.307). Scaling this effect by the standard deviation of GAI tone (0.581) implies an effect size of approximately 45.9%. In contrast, a one-standard-deviation increase in lagged GAI tone (0.581) results in only a 0.061 increase in human tone (0.105 × 0.581), which corresponds to about 19.9% of the standard deviation of human tone (0.307). These results suggest that most of the observed convergence in conversations is driven by GAI incorporating human input, rather than humans adjusting to the model. The evidence is consistent with the univariate patterns documented in Section 5.1 and provides support for H1.

We conduct two additional robustness checks to assess whether the baseline GMM results are robust to sample composition and alternative estimation methods. First, Table A1 shows that the coefficient of GAI's adaptation toward users remains concentrated between 0.863 and 0.866 across conversation-level resamples, whereas the human-side coefficient remains much smaller, ranging from 0.081 to 0.095. Second, Table A2 reports reduced-form OLS specifications that include lagged changes in tone. In the specification focusing on GAI's tone change, the coefficient on the human-to-GAI tone gap ranges from 0.546 to 0.995; they are substantially larger than those in the specification focusing on humans' tone change, i.e., 0.038 to 0.040. Although the OLS estimates may be subject to bias, the consistency of results across System GMM, Difference GMM, and OLS specifications with lagged controls strengthens our confidence in the inference.

5.3 Why GAI adapts to human users?

We next test H2-H4, uncovering the mechanisms underlying GAI adaptation to human users by analyzing cross-sectional variation in λ_{GAI}. Specifically, we estimate the following interacted dynamic models:

$$\begin{aligned}\widetilde{Tone}_{c,r}^{GAI} &= [1-(\lambda_0+\lambda_1 M_i)]\,\widetilde{Tone}_{c,r-1}^{GAI} + (\lambda_0+\lambda_1 M_i)\widetilde{Tone}_{c,r}^{Human} \\ &\quad + \gamma M_i + \varepsilon_{c,r}^{GAI},\end{aligned} \tag{5}$$

where M_i denotes the mechanism-specific moderator. The effective adjustment magnitude is $\lambda_0 + \lambda_1 M_i$, where λ_0 captures the baseline adjustment speed and λ_1 captures how the adjustment magnitude changes with the moderator. A positive λ_1 indicates larger convergence toward the counterpart when the moderator is higher. The coefficient captures the main effect of tone. We estimate these specifications using system GMM, as in the baseline dynamic specifications, to address persistence in tone and the resulting dynamic endogeneity.

Table 6 reports the analyses examining whether GAI adaptation becomes stronger when user inputs are more informative (H2a), GAI's internal knowledge is less precise (H2b), initial human–GAI signals are more independent (H3), and information-search tools are activated (H4). All panels estimate Equation (5) with various cross-sectional variables.

Panel A tests whether GAI adaptation varies with the quality of user-provided information. Column (1) uses $Tone^{Accurate}$ as the interaction term in Equation (5), while column (2) uses $FinLiteracy^{High}$ as the interaction term. $Tone^{Accurate}$ is an indicator equal to one if the sign of the user's initial tone is consistent with the direction of the subsequently realized firm fundamental, measured by the change in operating revenue for the forthcoming quarterly reports, and zero otherwise. This variable captures the accuracy of users' initial sentiment. $FinLiteracy^{High}$ is a user-level indicator equal to one if the proportion of finance-related words in the user's conversations in a given month exceeds the sample median, and zero otherwise.[9] This variable captures users' financial sophistication. As shown, the interaction coefficient is 1.834 in column (1), where $M_i = Tone^{Accurate}$, and 2.017 in column (2), where $M_i = FinLiteracy^{High}$. The interaction coefficient is positive and statistically significant at the 10% level for $Tone^{Accurate}$ and at the 1% level for $FinLiteracy^{High}$.[10] Therefore, we find the evidence supporting H2a: GAI's adaptation towards human users is selective, depending on the user's

[9]Finance-related words are identified using the finance dictionary provided by the Sogou Pinyin Input Method, available at https://pinyin.sogou.com/dict/detail/index/15127.

[10]The group-specific total adjustment speed is thus $\lambda_0+\lambda_1 M_i$. When this value exceeds one, it indicates overshooting: GAI adjusts beyond the change of humans, gradually closing the initial human–GAI tone gap.

informativeness.

Next, we examine whether GAI relies more on user inputs when its own domain knowledge is likely to be less precise. We report the results in Panel B. Because large language models are pre-trained on large-scale public text corpora, primarily including internet-based information (Brown et al., 2020; Ouyang et al., 2022), the amount of public firm-specific information available during the training period might shape the model's prior knowledge about a firm. We therefore use $Media^{Low}$ as the proxy for the lack of public information environments and, hence, weaker GAI domain knowledge. We define $Media^{Low}$ as a firm-level indicator equal to one if the number of original-title online media news mentions for the firm before the platform starts training is below the industry median, and zero otherwise. The interaction coefficient is 3.398 and statistically significant at the 10% level. This pattern is consistent with the idea that GAI places greater weight on user-provided information when its own prior knowledge is less comprehensive, supporting the interpretation that adaptation reflects selective information incorporation rather than biased sycophancy.

Furthermore, consistent with H3, Panel C examines whether information independency strengthens GAI's incentive to learn from users. Specifically, $Tone_{R0}^{GapHigh}$ is a conversation-level indicator equal to one if the initial-round tone gap between the human user and GAI in the conversation is above the median across conversations, and zero otherwise. The positive interaction coefficient ($\lambda_1 = 1.104$) indicates that GAI tone convergence is stronger when conversations begin with greater disagreement. This pattern is consistent with larger updating when user and GAI signals are less overlapping. As a caveat, the Hansen diagnostics suggest that this evidence should be interpreted cautiously.

Finally, we examine the role of search tools. H4 predicts that GAI adaptation is stronger when the model activates search-based retrieval, allowing it to retrieve external information and better interpret and validate the precision of user-provided inputs. Our data allow us to directly observe whether the model activates the search function in each conversation; specifically, we measure $WebSearch^{On}$ at the conversation level as an indi-

cator for whether the web search was used in the initial round of the conversation. Panel D presents the GAI equation with $M_i = WebSearch^{On}$. The interaction coefficient is positive and statistically significant at the 1% level ($\lambda_1 = 1.298$). These results indicate that search amplifies GAI-side adjustment, consistent with the prediction of H4. Although we cannot directly observe what GAI thinks, the observed increase in adaptation is consistent with the notion that retrieval improves the model's ability to interpret and verify user-provided information, leading to greater reliance on such inputs.

5.4 Outcomes of GAI adaptation

We then turn to H5, examining whether GAI adaptation improves both the GAI's and human user's information precision. Because the GMM requires sufficient observations, it is not feasible to directly estimate a conversation-specific λ_{GAI}. We therefore continue to rely on the cross-sectional partition approach to assess whether GAI-side adaptation is associated with the economic outcomes. Specifically, we estimate the following interacted dynamic specification:

$$\begin{aligned}\widetilde{Tone}_{c,r}^{GAI} &= [1-(\lambda_0+\lambda_1 O_i)]\,\widetilde{Tone}_{c,r-1}^{GAI} + (\lambda_0+\lambda_1 O_i)\widetilde{Tone}_{c,r}^{Human} \\ &\quad + \gamma O_i + \varepsilon_{c,r}^{GAI}.\end{aligned} \tag{6}$$

In Equation (6), O_i denotes the outcome-based moderator of interest, including user retention, the accuracy of the change in tones, or tone-signed future returns. Our underlying assumption is that if users repeatedly engage with the platform, tone accuracy improves, and the predictability of tone changes for future stock returns increases, then the updated beliefs become more informative. Consequently, this information enrichment yields superior economic outcomes. The effective GAI-side adjustment speed is $\lambda_0+\lambda_1 O_i$. Thus, λ_0 captures the baseline adjustment speed when $O_i = 0$, and λ_1 captures the incremental adjustment speed when $O_i = 1$. A positive λ_1 indicates that stronger GAI-side adaptation is associated with a higher value of the outcome variable. The coefficient γ captures the direct level effect of O_i on the GAI's tone, separate from its effect on the adjustment speed.

Table 7 reports the results. Panel A examines subsequent user engagement. If GAI adaptation indeed enhances the perceived value of the interaction, users should be more likely to continue using the platform. Column (1) considers $O_i = Comeback^{Human}$, an indicator equal to one if the user returns to ask firm-related financial questions afterward. The interaction coefficient is positive and significant ($\lambda_1 = 1.504$), suggesting that stronger GAI adaptation is associated with a higher likelihood of user return. To assess economic significance, we benchmark the magnitude against a 10% increase in GAI's adaptation relative to the baseline estimate of 0.868. The estimated interaction coefficient implies that such an increase is associated with a 5.8 percentage-point higher likelihood that the user returns to the platform and initiates a subsequent conversation.[11] This evidence is consistent with adaptive interactions enhancing user engagement, which is, to some extent, an intended objective of such systems.[12]

Panel B examines whether stronger GAI adaptation is associated with more accurate beliefs. If adaptation reflects efficient information extraction rather than mechanical sycophancy, stronger adaptation should be linked to reduced ex post belief distortion. Columns (1)–(2) consider two outcome-based measures: $Bias^{GAI}$ and $Bias^{Human}$. They capture whether the final belief becomes directionally less accurate, $Bias^{GAI}$ and $Bias^{Human}$ are indicators for deterioration in directional accuracy. For the GAI and the human user, respectively, the indicator equals one when the initial tone is directionally consistent with subsequently realized outcomes but the final tone is directionally inconsistent with those outcomes, and equals zero otherwise. The outcome are verifiable metric such as ROA, EPS, net profit, operating revenue, and so on, depending on the focus of conversation. The relevant outcome metric is manually identified and verified.

In addition, if our information-extraction mechanism is correct, stronger adaptation should increase the likelihood that GAI forms a more information-efficient belief, because adaptation allows GAI to extract and integrate information from both sides of the con-

[11] Let $\Delta\lambda = 0.1 \times 0.868 = 0.0868$ denote a 10% increase in GAI's adaptation relative to the baseline system-GMM estimate reported in Table 5. In Panel A, the interacted specification implies $\Delta\lambda = \lambda_1 \Delta Comeback^{Human}$. Therefore, $\Delta Comeback^{Human} = 0.0868/1.504 = 0.0577$, corresponding to approximately 5.8 percentage points.

[12] In private interviews with engineers from this and other large model providers, user retention is consistently described as a key target.

versation. We therefore construct $Informative^{GAI}$, a conversation-level indicator equal to one if the final-round GAI response is more accurate than the human user's final-round message relative to the realized outcome, and zero otherwise.

As shown, the results indicate that stronger GAI adaptation is associated with lower belief distortion and greater GAI informativeness. The interaction coefficients are negative for $Bias^{GAI}$ in column (1) ($\lambda_1 = -2.879$) and $Bias^{Human}$ in column (2) ($\lambda_1 = -1.202$), and positive for $Informative^{GAI}$ in column (3) ($\lambda_1 = 1.649$). These findings suggest that when GAI is more adaptive, both the human user and the GAI exhibit less bias, and the GAI's final-round response is more likely to be more accurate than the human user's final-round message.[13] Taken together, the evidence is consistent with GAI adaptation improving belief formation by selectively incorporating informative user inputs.

Last, capital markets, particularly firm-specific discussions, provide clues for constructing investment portfolios. We examine how the profit resulting from such portfolios is associated with GAI adaptation. If end-of-conversation sentiment reflects more accurate information, stronger adaptation should be associated with more profitable return signals for buy or sell positions formed based on the sentiment of the final tone. Tone-signed cumulative abnormal returns (CARs) are computed by first calculating daily abnormal returns using a market model, cumulating them over the relevant window (e.g., $[t+1, t+20]$ or $[t+2, t+20]$), and weighting by the final-round sentiment: neutral = 0, positive = 1, negative = -1. CAR^{Human} uses the human user's final tone, CAR^{GAI} uses the GAI's final tone.[14]

Panel C, Columns (1)–(4) report the regression results. Following prior studies (Kaniel

[13] Again, we benchmark the estimates against a 10% increase in GAI's adaptation relative to the baseline estimate of 0.868. Column (1) implies a 3.0 percentage-point lower probability that the GAI's final sentiment becomes less consistent with the subsequently realized outcome than its initial sentiment. Column (2) implies a 7.2 percentage-point lower probability that users revise their sentiment away from the direction subsequently supported by realized outcomes. Column (3) implies a 5.3 percentage-point higher probability that the final-round GAI response is more accurate than the human user's final-round message.

[14] For example, if a user expresses positive sentiment at the end of the conversation and the stock earns a post-conversation CAR of 10%, the tone-signed CAR equals $1 \times 10\% = 10\%$, corresponding to the return from buying the stock based on the user's final tone. Conversely, if the user expresses negative sentiment and the stock subsequently earns a CAR of 10%, the tone-signed CAR equals $-1 \times 10\% = -10\%$, reflecting the loss associated with taking a sell or short position based on the user's final tone.

et al., 2008; Kelley and Tetlock, 2013), we use a one-month event window and consider portfolio formation at the close on the conversation day, with returns measured from $t+1$ to $t+20$. We also report a more conservative specification that forms the portfolio at the next day's close, with returns measured from $t+2$ to $t+20$. Columns (1) and (3) use returns signed by the human user's final tone, whereas columns (2) and (4) use returns signed by the GAI's final tone. As shown, the interaction coefficients are positive and statistically significant across both windows and for both human- and GAI-signed returns, with estimates of 5.678 for $CAR^{Human}_{[1,20]}$, 5.871 for $CAR^{Human}_{[2,20]}$, 8.498 for $CAR^{GAI}_{[1,20]}$, and 9.423 for $CAR^{GAI}_{[2,20]}$. These results indicate that conversations with stronger GAI adaptation are followed by higher returns in the direction implied by the final tone of both the user and the GAI. However, the Hansen test fails in column (4), although the difference-in-Hansen tests do not reject in columns (2) and (4). We therefore caution the interpretation of Column (4). In terms of economic significance, a 10% increase in GAI's adaptation relative to the baseline estimate of 0.868 corresponds to about a 1.5% increase in profits based on the user's final-round sentiment and about a 1.0% increase in profits based on the GAI's final-round sentiment over the $[1, 20]$ window.[15] Overall, Panel C suggests that adaptive human–GAI interaction is linked not only to improved belief formation but also to economically meaningful investment return.

6 Additional Analysis of Mutual Adaptation

Having examined GAI adaptation alongside its underlying mechanisms and economic consequences, we now address one additional question: is GAI adaptation correlated with human adaptation? The answer to this final question sheds more light on why GAI's adaptation improves human users' information precision.

To address this, we leverage a bootstrap procedure to assess the correlation between

[15] With $\Delta\lambda = 0.0868$, the required change in the tone-signed return moderator is $\Delta CAR = \Delta\lambda/\lambda_1$. In column (1), $\Delta CAR^{Human}_{[1,20]} = 0.0868/5.678 = 0.0153$, or 1.5%. In column (2), $\Delta CAR^{GAI}_{[1,20]} = 0.0868/8.498 = 0.0102$, or 1.0%. Using the more conservative $[2, 20]$ window, the corresponding magnitudes are $0.0868/5.871 = 1.5\%$ for human-signed returns and $0.0868/9.423 = 0.9\%$ for GAI-signed returns. However, the results in Column (4) should be interpreted with caution because the Hansen test rejects the null hypothesis that the instruments are jointly exogenous and valid.

GAI-side and human-side adaptation. For each conversation size, we repeatedly resample the data 1,000 times and re-estimate the GMM adaptation coefficients. (Table A1 reports the results.) With this dataset, we examine the correlation between GAI's adaptation and human's adaptation. The estimated correlations are positive across all subsample sizes, ranging from approximately 0.18 to 0.25. This evidence implies that the bidirectional adaptation between both parties is positively correlated, suggesting that the GAI system and the human user engage in mutual adaptation to update their respective beliefs. Consequently, this collaborative updating process facilitates a more synchronized information exchange.

7 Conclusion

This study examines sentiment adaptation between human users and generative artificial intelligence (GAI) in financial markets. Motivated by a framework in which users and GAI possess substitutive information from different sources, we argue that adaptation in human–GAI interaction need not reflect mechanical sycophancy. Instead, it can arise endogenously through information extraction during multi-round conversations.

Using proprietary data on real-world firm-specific interactions from a leading Chinese GAI platform, we document three main findings. First, sentiment adaptation is asymmetric: GAI adjusts substantially to user sentiment, whereas user responses to GAI are much weaker. Second, this adaptation is selective rather than indiscriminate. GAI responds more strongly when user inputs are more informative, when GAI's domain knowledge is more limited, when their initial information is more independent to each other, and when search-based retrieval is activated. Third, these interaction dynamics have economically meaningful consequences. Stronger human–GAI adaptation is associated with more accurate belief updates and better trading performance, consistent with the view that adaptive interaction improves information extraction rather than merely reinforcing user bias.

Our findings contribute to the emerging literature on generative AI in capital markets. We move beyond treating AI outputs as static objects and instead study GAI as an

interactive intermediary whose responses evolve endogenously in conversation. A key insight emerging from our study is that GAI's adaptation towards human users appears to improve both its response quality and human users' information.

Several limitations remain. Our analysis focuses on one platform and one financial setting, and the exact magnitude of the economic effects may vary across models, users, and institutional environments. In addition, while our evidence is more consistent with information extraction than with pure sycophancy, future research can further disentangle these channels using settings with cleaner user identity or stronger identification. More generally, understanding when human–GAI interaction improves belief formation and when it amplifies bias remains an important agenda for research on GAI and capital markets.

References

Arellano, M. and Bond, S. (1991). Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *The Review of Economic Studies*, 58(2):277–297.

Baker, M. and Wurgler, J. (2007). Investor sentiment in the stock market. *Journal of economic perspectives*, 21(2):129–151.

Balasubramaniam, S. and Sabat, J. (2026). Ai as an informational intermediary. Available at SSRN 6075786.

Barberis, N. and Thaler, R. (2003). A survey of behavioral finance. In *Handbook of the Economics of Finance*, pages 1053–1128. Elsevier.

Bertomeu, J., Lin, Y., Liu, Y., and Ni, Z. (2025). The impact of generative AI on information processing: Evidence from the ban of chatgpt in italy. *Journal of Accounting and Economics*, page 101782.

Beyer, A., Cohen, D. A., Lys, T. Z., and Walther, B. R. (2010). The financial reporting environment: Review of the recent literature. *Journal of Accounting and Economics*, 50(2–3):296–343.

Blankespoor, E., Croom, J., and Grant, S. M. (2026a). Generative AI and investor processing of financial information. Available at SSRN.

Blankespoor, E., Dehaan, E., and Li, Q. (2026b). Generative AI in financial reporting. *Journal of Accounting Research*.

Blundell, R. and Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics*, 87(1):115–143.

Bradshaw, M. T., Ma, C., Yost, B. P., and Zou, Y. (2026). Generative AI use by capital market information intermediaries: Evidence from seeking alpha. *Journal of Accounting Research*.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. (2020). Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901.

Bushee, B. J., Gerakos, J., and Lee, L. F. (2018). Corporate jets and private meetings with investors. *Journal of Accounting and Economics*, 65(2–3):358–379.

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

Chandra, K., Kleiman-Weiner, M., Ragan-Kelley, J., and Tenenbaum, J. B. (2026). Sycophantic chatbots cause delusional spiraling, even in ideal bayesians. arXiv:2602.19141.

Chang, A. Y., Dong, X., Martin, X., and Zhou, C. (2026). AI democratization and trading inequality. *Journal of Accounting Research*.

Chen, S., Gao, M., Sasse, K., Hartvigsen, T., Anthony, B., Fan, L., Aerts, H., Gallifant, J., and Bitterman, D. S. (2025). When helpfulness backfires: LLMs and the risk of false medical information due to sycophantic behavior. *npj Digital Medicine*, 8:605.

Cheng, M., Lee, C., Khadpe, P., Yu, S., Han, D., and Jurafsky, D. (2026). Sycophantic AI decreases prosocial intentions and promotes dependence. *Science*, 391(6792):eaec8352.

Cheng, Q., Lin, P., and Zhao, Y. (2025a). Does generative AI facilitate investor trading? early evidence from chatgpt outages. *Journal of Accounting and Economics*, page 101821.

Cheng, Myraand Yu, S., Lee, C., Khadpe, P., Ibrahim, L., and Jurafsky, D. (2025b). Social sycophancy: A broader understanding of LLM sycophancy. arXiv preprint.

Coleman, B., Merkley, K., and Pacelli, J. (2022). Human versus machine: A comparison of robo-analyst and traditional research analyst investment recommendations. *The Accounting Review*, 97(5):221–244.

De Freitas, J., Agarwal, S., Schmitt, B., and Haslam, N. (2023). Psychological factors underlying attitudes toward AI tools. *Nature Human Behaviour*, 7(11):1845–1854.

Ecker, F., Li, X., Li, Y., and Wu, F. (2026). How stock market participants use generative artificial intelligence: Evidence from user-platform interaction data. *Journal of Accounting Research*.

Flannery, M. J. and Rangan, K. P. (2006). Partial adjustment toward target capital structures. *Journal of Financial Economics*, 79(3):469–506.

Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054.

Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review*, 35(4):519–530.

Healy, P. M. and Palepu, K. G. (2001). Information asymmetry, corporate disclosure, and the capital markets: A review of the empirical disclosure literature. *Journal of Accounting and Economics*, 31(1–3):405–440.

Huang, A. H., Wang, H., and Yang, Y. (2023). FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research*, 40(2):806–841.

Jussupow, E., Benbasat, I., and Heinzl, A. (2024). An integrative perspective on algorithm aversion and appreciation in decision-making. *MIS quarterly*, 48(4):1575–1590.

Kaniel, R., Saar, G., and Titman, S. (2008). Individual investor trading and stock returns. *The Journal of Finance*, 63(1):273–310.

Kelley, E. K. and Tetlock, P. C. (2013). How wise are crowds? insights from retail orders and stock returns. *The Journal of Finance*, 68(3):1229–1265.

Leib, M., Köbis, N., Rilke, R. M., Hagens, M., and Irlenbusch, B. (2024). Corrupted by algorithms? how AI-generated and human-written advice shape (dis) honesty. *The Economic Journal*, 134(658):766–784.

Li, E., Shen, M., Tu, Z., and Zhou, D. (2024). The promise and peril of generative AI: Evidence from gpt as sell-side analysts. *arXiv preprint arXiv:2412.01069*.

Loughran, T. and McDonald, B. (2011). When is a liability not a liability? textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, 66(1):35–65.

Naddaf, M. (2025). AI chatbots are sycophants — researchers say it's harming science. *Nature*, 647:13–14. News.

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica*, 49(6):1417–1426.

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., and Lowe, R. (2022). Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744.

Perez, E., Ringer, S., Lukosiute, K., Nguyen, K., Chen, E., Heiner, S., Pettit, C., Olsson, C., Kundu, S., Kadavath, S., et al. (2023). Discovering language model behaviors with model-written evaluations. arXiv:2212.09251.

Rabin, M. and Schrag, J. L. (1999). First impressions matter: A model of confirmatory bias. *Quarterly Journal of Economics*, 114(1):37–82.

Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askell, A., Bowman, S., Durmus, E., Hatfield-Dodds, Z., Johnston, S., Kravec, S., et al. (2024). Towards understanding sycophancy in language models.

Sheng, J., Sun, Z., Yang, B., and Zhang, A. L. (2026). Generative AI and asset management. *Review of Financial Studies*.

Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of Finance*, 62(3):1139–1168.

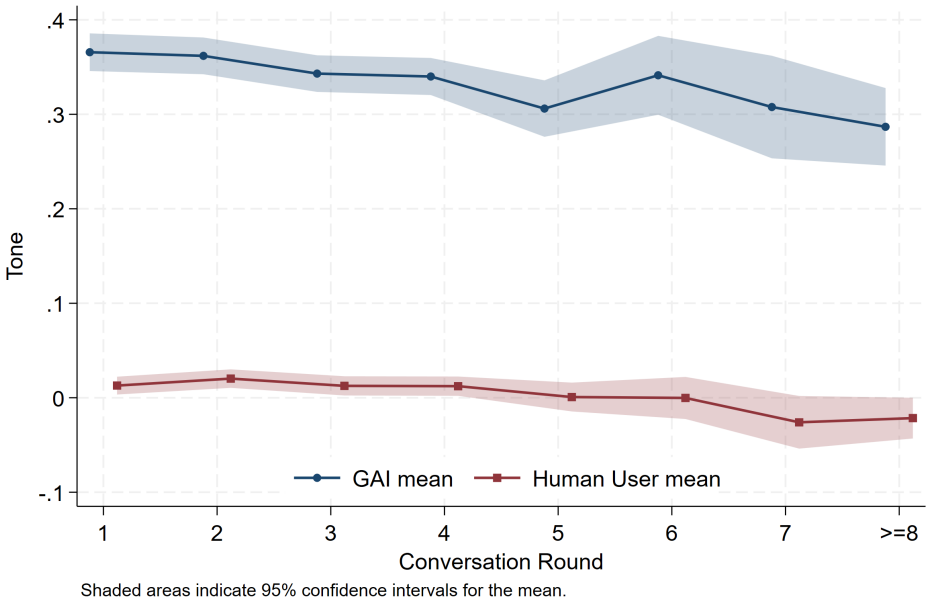
Wintoki, M. B., Linck, J. S., and Netter, J. M. (2012). Endogeneity and the dynamics of internal corporate governance. *Journal of Financial Economics*, 105(3):581–606.

Tables and Figures

Figure 1: Sentiment dynamics over conversation rounds

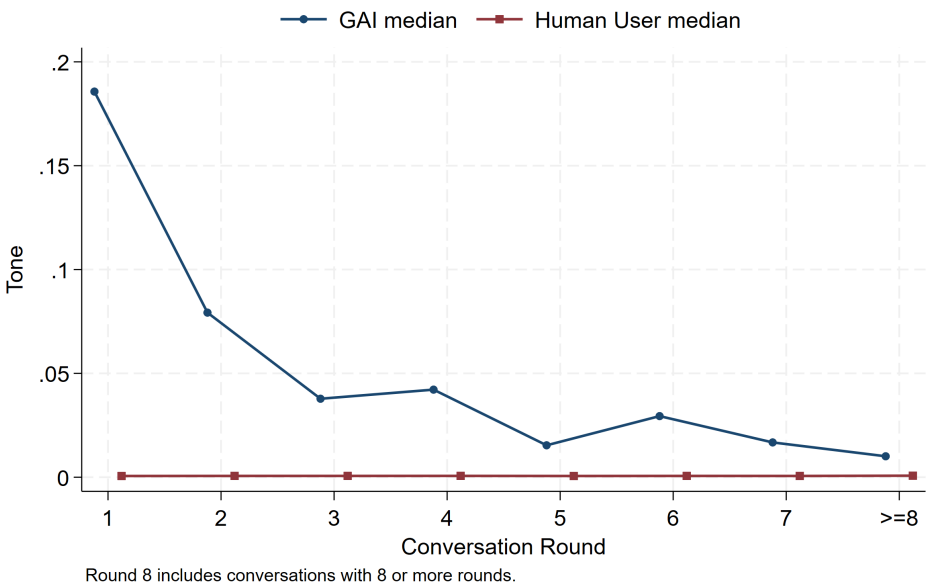
This figure plots the dynamics of GAI and human user tone over conversation rounds. Panel A reports the mean tone in each round, where the blue line represents GAI tone and the red line represents human user tone; shaded areas indicate 95% confidence intervals around the mean. Panel B reports the median tone in each round. The horizontal axis indicates the sequence of rounds within a conversation, where >=8 presents round 8 or above. The vertical axis reports the tone measure constructed from the finance-domain sentiment classifier.

Panel A: Mean tone



Shaded areas indicate 95% confidence intervals for the mean.

Panel B: Median tone



Round 8 includes conversations with 8 or more rounds.

Table 1: Sample selection of human–GAI conversations

This table reports how the raw platform logs are sequentially filtered into the final sample of firm-related human–GAI conversations. Each row shows the remaining share of conversation-rounds, distinct conversations, and unique firms after applying the corresponding filter. Filters remove duplicate firm matches, error records, multimodal interactions, missing or abnormal text, text-processing tasks (e.g., translation, polishing, rewriting, content expansion), single-round conversations, observations without firm-level controls, and dialogues classified as non–capital-market content by a fine-tuned Chinese BERT classifier.

	Sample selection	**Share of conversation rounds**	**Share of conversations**	**Share of firms**
(1)	Matched conversations with listed firms from Sep–Dec 2024 on the platform.	100.0%	100.0%	100.0%
(2)	Exclude conversations containing duplicate firms.	76.0%	79.1%	92.7%
(3)	Exclude conversations containing any observations with `status == "error"`.	61.6%	67.7%	91.6%
(4)	Exclude groups involving multimodal interactions.	54.4%	61.1%	90.6%
(5)	Exclude observations where prompt or answer text is missing, empty, whitespace-only, or with prompt length ≤ 4 or > 1000 characters.	46.0%	52.9%	89.8%
(6)	Exclude conversations related to text-processing tasks (translation, polishing, rewriting, content expansion).	9.2%	13.0%	71.6%
(7)	Exclude conversations containing only a single round.	8.6%	10.6%	66.6%
(8)	Exclude conversations linked to firms lacking control variables.	8.3%	10.2%	63.9%
(9)	Employ a fine-tuned Chinese-based BERT classifier to remove non–capital-market dialogues.	4.5%	5.9%	60.0%

Table 2: Sample Firms versus Full Population of Chinese Listed Firms

This table compares the average firm characteristics of the sample firms used in our analysis with those of the full population of Chinese listed firms at the end of 2024. The sample firms are firms included in our final analysis sample after applying all filters (N = 2,266). The full listed firms include all Chinese listed firms with available firm-level characteristics at the end of 2024. The difference is calculated as the mean value for sample firms minus the mean value for the full population of Chinese listed firms. *Log Assets* is the natural logarithm of total assets. *Leverage* is total liabilities divided by total assets. *Return on Assets* is net income divided by total assets. *Book-to-Market Ratio* is the book value of equity divided by the market value of equity. *Log Listed month* is the natural logarithm of one plus the number of months since the firm's listing date. *Stock Return* is the firm's average stock return over the last five trading days of 2024. *Analyst Coverage* is the number of analysts following the firm in 2024. Statistical significance of the differences is based on two-sample mean-comparison tests. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Sample firms (N = 2,266)	Full listed firms (N = 5,549)	Difference (=Sample−Full)
Log Assets	22.817	22.014	0.802***
Leverage	0.440	0.396	0.044***
Return on Assets	0.018	0.020	−0.001
Book-to-Market Ratio	0.622	0.584	0.038*
Log Listed Month	4.842	4.388	0.454***
Stock Return	0.004	0.006	−0.002
Analyst Coverage	22.239	10.920	11.319***

Table 3: Descriptive statistics

This table reports descriptive statistics for the variables used in the empirical analysis. Panel A reports demeaned tone variables in the sequential-round sample used for the dynamic adjustment estimation, and Panel B reports cross-sectional moderator variables used in the mechanism analyses. Panel C reports the outcome variables used to examine the economic consequences of GAI adaptation. The table reports the number of observations, the mean, the standard deviation, and the 10th, 50th, and 90th percentiles. See Appendix 8 for detailed variable definitions.

Variable	N	Mean	SD	p10	p50	p90
Panel A: Multi-round tone measures						
$\widetilde{Tone}^{GAI}$	8,105	−0.025	0.581	−0.705	−0.074	0.663
$\widetilde{Tone}^{Human}$	8,105	−0.012	0.307	−0.159	−0.009	0.116
Panel B: Mechanism variables						
$Tone^{Accurate}$	8,105	0.547	0.498	0	1	1
$FinLiteracy^{High}$	8,105	0.459	0.498	0	0	1
$Media^{Low}$	8,105	0.259	0.438	0	0	1
$WebSearch^{On}$	8,105	0.791	0.407	0	1	1
$Tone_{R0}^{GapHigh}$	8,105	0.573	0.495	0	1	1
Panel C: Outcome variables						
$Comeback^{Human}$	8,105	0.742	0.438	0	1	1
$Bias^{GAI}$	8,105	0.156	0.363	0	0	1
$Bias^{Human}$	8,105	0.160	0.367	0	0	1
$Informative^{GAI}$	8,105	0.271	0.445	0	0	1
$\text{CAR}_{[1,20]}^{Human}$	8,105	0.000	0.038	0	0	0
$\text{CAR}_{[2,20]}^{Human}$	8,105	0.000	0.035	0	0	0
$\text{CAR}_{[1,20]}^{GAI}$	8,105	−0.005	0.086	−0.081	0	0.040
$\text{CAR}_{[2,20]}^{GAI}$	8,105	−0.005	0.082	−0.078	0	0.039

Table 4: Sentiment dynamics from the initial round to the final round

This table compares tone in the initial round (R_0) and the final round (R^{end}) of the same multi-round conversation. $Tone^{GAI}$ denotes the tone of the GAI response, and $Tone^{Human}$ denotes the tone of the human user's message. Tone is measured as the predicted positive-tone probability minus the predicted negative-tone probability from the finance-domain sentiment classifier. The mean difference is defined as $R^{end} - R_0$ within a conversation, and statistical significance is based on a paired t-test of whether the mean within-conversation difference equals zero. The median difference is also defined as $R^{end} - R_0$ within a conversation, and statistical significance is based on a Wilcoxon signed-rank test of whether the median within-conversation difference equals zero. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	R_0		R^{end}		$R^{end} - R_0$	
	Mean	Median	Mean	Median	Mean diff.	Median diff.
$Tone^{GAI}$	0.381	0.168	0.360	0.048	-0.021^{***}	-0.120^{***}
$Tone^{Human}$	0.023	0.001	0.032	0.001	0.008^{***}	2.50×10^{-5}

Table 5: Human–GAI tone adaptation in sequential rounds

This table reports baseline estimates of tone adjustment in sequential rounds using dynamic panel GMM. Columns (1)–(2) estimate the GAI adjustment based on Equation (2), and columns (3)–(4) estimate the human user adjustment based on Equation (4). λ_{GAI} and λ_{Human} denote the adjustment speeds of GAI and the human user, respectively; a larger coefficient implies faster adaptation toward the counterpart's tone, whereas a smaller coefficient indicates greater persistence in one's own prior tone. Difference GMM uses the differenced specification only, whereas system GMM jointly estimates the level and differenced equations. The instrument sets differ across the two equations. For the GAI equation in columns (1)–(2), the differenced equation uses lagged levels of GAI tone dated $r-2$ and $r-3$, while the level equation in the system GMM specification uses lagged differences of both GAI tone and human tone dated $r-2$. For the human-user equation in columns (3)–(4), the differenced equation uses lagged levels of GAI tone dated $r-2$ and $r-3$ together with lagged human tone dated $r-4$, while the level equation in the system GMM specification uses lagged differences of both GAI tone and human tone dated $r-3$. AR(1) and AR(2) are Arellano–Bond tests for first- and second-order serial correlation in the residuals, under the null of no serial correlation. Hansen and Diff-in-Hansen are tests of overidentifying restrictions, where the null hypothesis is that the instruments are valid. Standard errors are clustered by conversation and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	GAI		Human	
	(1) Difference GMM	(2) System GMM	(3) Difference GMM	(4) System GMM
λ_{GAI}	0.865***	0.868***		
	(0.024)	(0.023)		
λ_{Human}			0.107***	0.105***
			(0.021)	(0.020)
AR(1) p-value	0.000	0.000	0.000	0.000
AR(2) p-value	0.820	0.798	0.898	0.898
Hansen p-value	0.413	0.289	0.657	0.935
Diff-in-Hansen p-value		0.220		0.976

Table 6: Mechanism analysis

This table reports mechanism tests based on the interacted system generalized model in Equation (5), focusing on GAI adjustment toward user tone. The estimated model is $\widetilde{Tone}_{c,r}^{GAI} = [1-(\lambda_0+\lambda_1 M_i)]\widetilde{Tone}_{c,r-1}^{GAI} + (\lambda_0+\lambda_1 M_i)\widetilde{Tone}_{c,r}^{Human} + \gamma M_i + \varepsilon_{c,r}^{GAI}$, where M_i denotes a mechanism-specific moderator. The coefficient λ_0 captures the baseline adjustment speed when $M_i = 0$, and λ_1 captures the incremental adjustment speed when $M_i = 1$. The coefficient γ captures the direct level effect of the moderator M_i. Panel A examines user information quality. $Tone^{Accurate}$ is an indicator equal to one if the sign of the user's initial tone is consistent with the direction of the subsequently realized firm fundamental, measured by the change in operating revenue for the given quarter, and zero otherwise. $FinLiteracy^{High}$ is an indicator equal to one if the user's finance-word ratio is above the sample median in a given month, and zero otherwise. Panel B examines GAI domain knowledge. $Media^{Low}$ is a firm-level indicator equal to one if the number of original-title online media news mentions for the firm before the platform starts training is below the industry median, and zero otherwise. Standard errors are clustered at the conversation level and reported in parentheses. Arellano–Bond AR(1) and AR(2) tests assess first- and second-order serial correlation in the differenced residuals. Hansen and Difference-in-Hansen tests assess the validity of the overidentifying restrictions and the additional moment conditions used in the system generalized method of moments estimator. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. User information quality

M_i	(1) $Tone^{Accurate}$	(2) $FinLiteracy^{High}$
λ_1	1.834*	2.017***
	(0.987)	(0.613)
λ_0	−0.112	−0.031
	(0.527)	(0.273)
γ	−0.015	0.009
	(0.016)	(0.021)
AR(1) p-value	0.000	0.000
AR(2) p-value	0.796	0.951
Hansen p-value	0.555	0.626
Diff-in-Hansen p-value	0.372	0.465

Panel B. GAI domain knowledge

M_i	(1) $Media^{Low}$
λ_1	3.398*
	(1.900)
λ_0	−0.069
	(0.519)
γ	−0.029
	(0.044)
AR(1) p-value	0.000
AR(2) p-value	0.195
Hansen p-value	0.572
Diff-in-Hansen p-value	0.396

Table 6: Mechanism analysis (continued)

This continued table reports additional mechanism tests based on the interacted system generalized model in Equation (5), focusing on GAI adjustment toward user tone. The estimated model is $\widetilde{Tone}_{c,r}^{GAI} = [1-(\lambda_0+\lambda_1 M_i)]\widetilde{Tone}_{c,r-1}^{GAI} + (\lambda_0+\lambda_1 M_i)\widetilde{Tone}_{c,r}^{Human} + \gamma M_i + \varepsilon_{c,r}^{GAI}$, where M_i denotes a mechanism-specific moderator. The coefficient λ_0 captures the baseline adjustment speed when $M_i = 0$, and λ_1 captures the incremental adjustment speed when $M_i = 1$. Thus, the effective adjustment speed is $\lambda_0 + \lambda_1 M_i$. The coefficient γ captures the direct level effect of the moderator M_i. Panel C examines information complementarity in the initial round. $Tone_{R0}^{GapHigh}$ is a conversation-level indicator equal to one if the initial-round tone gap between the human user and GAI in the conversation is above the median across conversations, and zero otherwise. Panel D examines the role of internet search tools. $WebSearch^{On}$ is an indicator equal to one if internet search is activated in the initial round of the conversation, and zero otherwise. Standard errors are clustered at the conversation level and reported in parentheses. Arellano–Bond AR(1) and AR(2) tests assess first- and second-order serial correlation in the differenced residuals. Hansen and Difference-in-Hansen tests assess the validity of the overidentifying restrictions and the additional moment conditions used in the system generalized method of moments estimator. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel C. Information complementarity

M_i	(1) $Tone_{R0}^{GapHigh}$
λ_1	1.104***
	(0.138)
λ_0	0.291***
	(0.066)
γ	0.075***
	(0.015)
AR(1) p-value	0.000
AR(2) p-value	0.901
Hansen p-value	0.000
Diff-in-Hansen p-value	0.000

Panel D. Internet search tools

M_i	(1) $WebSearch^{On}$
λ_1	1.298***
	(0.246)
λ_0	−0.163
	(0.195)
γ	0.004
	(0.009)
AR(1) p-value	0.000
AR(2) p-value	0.888
Hansen p-value	0.508
Diff-in-Hansen p-value	0.354

Table 7: Outcome Analysis

This table reports the outcome analysis based on the interacted dynamic GAI tone-adjustment in Equation (6). The estimated model is $\widetilde{Tone}_{c,r}^{GAI} = [1 - (\lambda_0 + \lambda_1 O_i)]\widetilde{Tone}_{c,r-1}^{GAI} + (\lambda_0 + \lambda_1 O_i)\widetilde{Tone}_{c,r}^{Human} + \gamma O_i + \varepsilon_{c,r}^{GAI}$, where O_i denotes an outcome-based moderator. The coefficient λ_0 captures the baseline GAI adjustment speed when $O_i = 0$, and λ_1 captures the incremental adjustment speed associated with the moderator O_i. The coefficient γ captures the direct level effect of the moderator. Panel A examines user retention. $Comeback^{Human}$ is an indicator equal to one if the same user continues to ask about the firm's financial information during the sample period, and zero otherwise. Panel B examines whether stronger GAI adjustment is associated with belief distortion and GAl informativeness. $Bias^{GAI}$ is an indicator equal to one if the initial GAI sentiment is directionally consistent with the subsequently realized outcome but the final-round GAI sentiment is not, and zero otherwise. $Bias^{Human}$ is defined analogously using the human user's sentiment. $Informative^{GAI}$ is an indicator equal to one if the final-round GAI signal is more accurate than the human user's signal, and zero otherwise. Standard errors are clustered at the conversation level and reported in parentheses. Arellano–Bond AR(1) and AR(2) tests assess first- and second-order serial correlation in the differenced residuals. Hansen and Difference-in-Hansen tests assess the validity of the overidentifying restrictions and the additional moment conditions used in the system generalized method of moments estimator. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. User retention

O_i	(1) $Comeback^{Human}$
λ_1	1.504**
	(0.694)
λ_0	−0.248
	(0.516)
γ	−0.010
	(0.011)
AR(1) p-value	0.000
AR(2) p-value	0.655
Hansen p-value	0.517
Diff-in-Hansen p-value	0.441

Panel B. Bias and GAI informativeness

O_i	(1) $Bias^{GAI}$	(2) $Bias^{Human}$	(3) $Informative^{GAI}$
λ_1	−2.879*	−1.202**	1.649*
	(1.485)	(0.579)	(0.959)
λ_0	1.404***	1.097***	−0.003
	(0.282)	(0.116)	(0.507)
γ	−0.054	−0.042*	−0.012
	(0.038)	(0.023)	(0.015)
AR(1) p-value	0.000	0.000	0.000
AR(2) p-value	0.625	0.634	0.800
Hansen p-value	0.660	0.144	0.369
Diff-in-Hansen p-value	0.611	0.196	0.315

Table 7: Outcome Analysis (continued)

This continued table reports the outcome analysis based on the interacted dynamic GAI tone-adjustment in Equation (6). The estimated model is $\widetilde{Tone}_{c,r}^{GAI} = [1-(\lambda_0+\lambda_1 O_i)]\widetilde{Tone}_{c,r-1}^{GAI} + (\lambda_0 + \lambda_1 O_i)\widetilde{Tone}_{c,r}^{Human} + \gamma O_i + \varepsilon_{c,r}^{GAI}$, where O_i denotes an outcome-based moderator. The coefficient λ_0 captures the baseline GAI adjustment speed when $O_i = 0$, and λ_1 captures the incremental adjustment speed associated with the moderator O_i. The coefficient γ captures the direct level effect of the moderator. In Panel C, O_i is a tone-signed cumulative abnormal return over the future return window indicated in each column. CAR^{Human} is constructed by multiplying cumulative abnormal returns by the sign of the human user's final-round tone, where positive tone is coded as 1, neutral tone as 0, and negative tone as -1. CAR^{GAI} is constructed analogously using the sign of the GAI final-round tone. Thus, these variables capture the future abnormal return implied by the end-of-conversation tone signal. $CAR_{[1,20]}$ and $CAR_{[2,20]}$ denote cumulative abnormal returns over trading days $[t+1, t+20]$ and $[t+2, t+20]$ following the conversation date, respectively. Standard errors are clustered at the conversation level and reported in parentheses. Arellano–Bond AR(1) and AR(2) tests assess first- and second-order serial correlation in the differenced residuals. Hansen and Difference-in-Hansen tests assess the validity of the overidentifying restrictions and the additional moment conditions used in the system generalized method of moments estimator. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel C. Tone-signed return

O_i	(1) $CAR^{Human}_{[1,20]}$	(2) $CAR^{GAI}_{[1,20]}$	(3) $CAR^{Human}_{[2,20]}$	(4) $CAR^{GAI}_{[2,20]}$
λ_1	5.678**	8.498**	5.871**	9.423**
	(2.731)	(3.613)	(2.879)	(4.753)
λ_0	0.859***	0.937***	0.861***	0.941***
	(0.025)	(0.046)	(0.024)	(0.053)
γ	0.356	0.491	0.320	0.519
	(0.401)	(0.360)	(0.412)	(0.431)
AR(1) p-value	0.000	0.000	0.000	0.000
AR(2) p-value	0.731	0.261	0.789	0.244
Hansen p-value	0.732	0.141	0.734	0.029
Diff-in-Hansen p-value	0.934	1.000	0.904	1.000

8 Variable Definitions

This appendix defines the main variables used in the empirical analysis.

Variable	Definition
Tone Measures	
$Tone_0^{GAI}$, $Tone_0^{Human}$	The tone of the GAI response and the user prompt, respectively, in the initial round of a conversation. Tone is measured using a finance-domain sentiment classifier, with higher values indicating more positive tone.
$\widetilde{Tone}^{GAI}$, $\widetilde{Tone}^{Human}$	Demeaned GAI tone and user tone, respectively, after removing firm and date fixed effects. These residualized tone measures capture within-firm-date variation and are used in the dynamic adjustment analysis.
Dynamic Adjustment and Estimation Notation	
λ_{GAI}, λ_{Human}	The speed of tone adjustment for the GAI and the user, respectively. These coefficients capture how strongly one side adjusts its tone toward the other side in the dynamic specification.
λ_0, λ_1, γ	Coefficients from interacted specifications. In the mechanism analyses, the effective adjustment speed is $\lambda_0 + \lambda_1 M_i$; in the outcome analyses, it is $\lambda_0 + \lambda_1 O_i$. The coefficient γ captures the corresponding level effect.
M_i	A mechanism-specific moderator used to examine whether adjustment intensity varies across information environments, user sophistication, or other cross-sectional settings.
O_i	An outcome-based moderator used to examine whether adjustment intensity is associated with user retention, belief accuracy, or tone-signed future returns.
AR(1), AR(2)	Arellano–Bond tests for first-order and second-order serial correlation in the differenced residuals.
Hansen, Diff-in-Hansen	Tests of overidentifying restrictions and the validity of additional moment conditions in GMM estimation.
Mechanism Variables	
$Tone_{R0}^{GapHigh}$	A conversation-level indicator equal to one if the initial-round tone gap between the human user and GAI in the conversation is above the median across conversations, and zero otherwise. It is used to examine whether adjustment is stronger when the initial divergence is larger.
$Media^{Low}$	A firm-level indicator equal to one if the number of original-title online media news mentions for the firm is below the industry median before the platform starts training, and zero otherwise.
$WebSearch^{On}$	A conversation-level indicator equal to one if internet search is activated in the initial round of the conversation, and zero otherwise.
$Tone^{Accurate}$	A conversation-level indicator equal to one if the sign of the user's initial tone is consistent with the direction of the subsequently realized fundamental, measured by the change in operating revenue for the given quarter, and zero otherwise.
$FinLiteracy^{High}$	A user-level indicator equal to one if the user's finance-word ratio is above the sample median in a given month, and zero otherwise.

Continued on next page

Table *continued*

Variable	Definition
Outcome Variables	
$Comeback^{Human}$	An indicator equal to one if the user returns to the platform for a subsequent interaction after the focal conversation, and zero otherwise.
$Bias^{GAI}$, $Bias^{Human}$	Indicators for whether the GAI's or the user's tone, respectively, shifts from being directionally consistent with realized outcomes at the beginning of the conversation to being inconsistent by the end of the conversation.
$Informative^{GAI}$	A conversation-level indicator equal to one if the GAI response in the last round is more accurate than the human user's final-round message, and zero otherwise.
$CAR^{Human}_{[1,20]}$, $CAR^{Human}_{[2,20]}$	Tone-signed cumulative abnormal returns based on the final-round human user tone, measured over trading days $[t+1, t+20]$ or $[t+2, t+20]$ after the conversation date; positive (negative) tone is assigned $+1$ (-1), and neutral tone is assigned 0.
$CAR^{GAI}_{[1,20]}$, $CAR^{GAI}_{[2,20]}$	Tone-signed cumulative abnormal returns based on the final-round GAI tone, measured over trading days $[t+1, t+20]$ or $[t+2, t+20]$ after the conversation date; positive (negative) tone is assigned $+1$ (-1), and neutral tone is assigned 0.

9 Examples of human–GAI Interactions

This appendix presents representative examples of human–GAI interactions. All firm names and identifiable information are anonymized. These examples illustrate how GAI responses evolve during multi-round conversations and how adaptation to user opinions may emerge.

Example 1

> **User (Round 1):** It is reported that orders for [Firm A]'s new model surged during the National Day holiday, and deliveries have been extremely busy. Does this imply strong market demand?
>
> **GAI Assistant (Round 1):** According to public statements by senior management of [Firm A], orders for its new model increased significantly during the holiday period, and delivery operations have been running at high capacity. This suggests strong consumer demand and growing popularity of the model in the market.
>
> Furthermore, it has been mentioned that orders within the first 24 hours after launch exceeded the monthly sales of comparable models in the same segment, indicating strong market performance and high consumer interest.
>
> Overall, these signals reflect [Firm A]'s strong momentum in the new energy vehicle market and its increasing competitiveness.
>
> **User (Round 2):** It is also said that orders for the model exceeded the monthly sales of comparable luxury sedans within 24 hours of launch. Yet [Firm A] itself has not publicized this aggressively. Is the firm being too low-profile despite its market success?
>
> **GAI Assistant (Round 2):** While management did mention strong early demand, the firm has not officially disclosed detailed order figures, possibly due to a preference for a more conservative communication strategy. This suggests that [Firm A] may intentionally adopt a low-profile approach in its marketing.
>
> **User:** I actually think [Firm A] is being too conservative. In such a competitive market, being overly low-profile might cause it to fall behind more aggressive competitors.
>
> **GAI Assistant (Round 3):** Your perspective is very reasonable. A conservative and low-profile strategy can indeed be interpreted as a sign of confidence and long-term orientation.
>
> There are several possible explanations: (1) a focus on product quality and technological improvement, (2) avoiding excessive marketing that may create unrealistic expectations, (3) positioning the firm as a technology-driven rather than marketing-driven brand, and (4) emphasizing long-term stability over short-term hype.
>
> Such a strategy may allow [Firm A] to maintain sustainable growth and competitive advantage.

Example 2

User (Round 1): [Firm B]'s global sales in September look strong. Does this indicate continued growth?

GAI Assistant (Round 1): [Firm B]'s sales performance has been strong, with total sales reaching over 120,000 units and significant year-over-year growth. New energy vehicle sales grew particularly rapidly, indicating strong momentum.

User (Round 2): But [Firm B] does not seem to pursue aggressive expansion. Is that sustainable?

GAI Assistant (Round 2): A non-aggressive expansion strategy may reflect a focus on sustainable development. Key advantages include better risk control, stronger quality assurance, improved financial stability, and more flexibility in adapting to market changes.

User (Round 3): I think expansion speed is less important than quality and profitability.

GAI Assistant (Round 3): Your view is very comprehensive and insightful. Sales growth is only one dimension of performance; product quality, brand reputation, and profitability are equally important.

Focusing on quality and sustainable profitability can help firms build long-term competitiveness and stand out in a highly competitive market.

Example 3

User (Round 1): [Firm C]'s high leverage may indicate financial risk. Is this concerning?

GAI Assistant (Round 1): [Firm C] has experienced significant financial challenges, including accumulated losses, declining revenues, and liquidity pressure. These factors may increase financial risk and affect investor confidence.

User (Round 2): If fundamentals are weak, why would institutional investors still hold the stock?

GAI Assistant (Round 2): Institutional investors may maintain positions due to expectations of future recovery, improvements in corporate governance, or long-term strategic considerations.

User (Round 3): Where do state-owned investors come from?

GAI Assistant (Round 3): State-owned investors typically originate from government-backed investment entities. These investors may acquire shares through negotiated transactions and aim to support long-term corporate restructuring.

10 Robustness Tests

Table A1: Mutual adaptation

This table reports bootstrap evidence on the stability of the GMM adaptation estimates across subsamples of different conversation sizes. For each listed conversation size, conversations are repeatedly resampled 1,000 times and the adjustment coefficients are re-estimated. Panel A reports the bootstrap mean and the 95% confidence interval for λ_{GAI} and λ_{Human}. Panel B reports the mean bootstrap correlation between λ_{GAI} and λ_{Human} together with its 95% confidence interval.

Panel A. Adaptation coefficients across conversation sizes				
	λ_{GAI}		λ_{Human}	
# Conversation	Mean	95% CI	Mean	95% CI
1000	0.864	[0.734, 1.007]	0.081	[−0.033, 0.195]
2000	0.864	[0.777, 0.957]	0.085	[0.006, 0.168]
3000	0.866	[0.797, 0.936]	0.092	[0.019, 0.158]
4000	0.863	[0.807, 0.929]	0.092	[0.033, 0.145]
5000	0.865	[0.817, 0.916]	0.095	[0.048, 0.141]
Panel B. Correlation between λ_{GAI} estimates (bootstrap)				
# Conversation	Mean correlation	95% CI		
1000	0.225	[0.155, 0.294]		
2000	0.180	[0.115, 0.245]		
3000	0.203	[0.139, 0.267]		
4000	0.247	[0.183, 0.311]		
5000	0.205	[0.139, 0.270]		

Table A2: OLS Robustness Test with Lagged Controls

This table reports a robustness check using OLS specifications with lagged changes in tone as controls. Columns (1)–(3) use $\Delta Tone^{GAI}$ as the dependent variable, and columns (4)–(6) use $\Delta Tone^{Human}$ as the dependent variable. Columns (2) and (5) include firm and date fixed effects, columns (3) and (6) include conversation fixed effects, and columns (1) and (4) are pooled specifications without fixed effects. The key regressor is the tone gap shown in the first row of each block, and $L.$ and $L2.$ denote the first and second lags of tone changes. Standard errors are clustered by conversation and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	GAI			Human		
	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Variable:		$\Delta Tone^{GAI}$			$\Delta Tone^{Human}$	
$Tone_r^{Human} - Tone_{r-1}^{GAI}$	0.546***	0.655***	0.995***			
	(0.015)	(0.019)	(0.031)			
$Tone_{r-1}^{GAI} - Tone_{r-1}^{Human}$				0.038***	0.038***	0.040***
				(0.006)	(0.008)	(0.012)
Log Assets		0.645			0.077	
		(0.528)			(0.264)	
Leverage		−0.817			−0.133	
		(0.910)			(0.609)	
Return on Assets		−1.862			1.432	
		(1.770)			(0.919)	
Book-to-Market Ratio		−0.014			0.005	
		(0.056)			(0.026)	
Log Listed Month		−2.867**			−0.222	
		(1.411)			(0.913)	
Stock Return		0.492			0.245	
		(0.369)			(0.191)	
Analyst Coverage		0.001			−0.001	
		(0.002)			(0.001)	
$L.\Delta Tone^{GAI}$	−0.252***	−0.221***	−0.096***			
	(0.015)	(0.017)	(0.027)			
$L2.\Delta Tone^{GAI}$	−0.125***	−0.115***	−0.071***			
	(0.012)	(0.014)	(0.017)			
$L.\Delta Tone^{Human}$				−0.560***	−0.602***	−0.715***
				(0.017)	(0.020)	(0.027)
$L2.\Delta Tone^{Human}$				−0.313***	−0.332***	−0.398***
				(0.018)	(0.019)	(0.024)
Observations	8,105	8,029	8,105	8,105	8,029	8,105
R^2	0.4098	0.5008	0.6453	0.2620	0.3493	0.4510
Firm FE	No	Yes	No	No	Yes	No
Date FE	No	Yes	No	No	Yes	No
Conversation FE	No	No	Yes	No	No	Yes