# Discussion of *Does Artificial Intelligence Reduce Investor Disagreement? Evidence from AI–Human Interactions Data*

Huang, Shao, Song, Wang, Xue (2026)

Discussant: Danqing Mei

# What the paper does

**Research question.** Does GenAI usage make investor beliefs converge or diverge?

**Data.** Proprietary multi-round user–chatbot conversation logs from a major Chinese GenAI platform (Sep–Dec 2024), matched to listed firms.

**Measurement.**

- Tone (FinBERT) of each user prompt and AI response.
- Disagreement = cross-user variance of tone at firm–date–round level.
- Decomposed into **within-topic** and **between-topic** components.

**Extension.** U.S. diff-in-diff around ChatGPT launch (Nov 2022), looking at analyst dispersion, turnover, IVOL, spreads, DO/IA around earnings announcements.

# Main findings

1. User tone dispersion **rises** across conversation rounds (+11% R1$\rightarrow$R2; $|\text{tone}|$ +25%).
2. The increase is driven **almost entirely by between-topic dispersion**; within-topic is flat.
3. Users update toward AI tone (Table 6).
4. AI tone predicts future returns; human tone does not $\Rightarrow$ "information, not noise."
5. Strong **topic stickiness**; additional rounds add no incremental predictability $\Rightarrow$ called "behavioral."
6. Post-ChatGPT: more dispersion at earnings announcements, but *less* information asymmetry.

# The big idea — and why it's interesting

## Personalized precision without belief convergence.

AI improves the *signal within* the topic each user cares about, but users persistently care about *different* topics. Aggregate disagreement can rise even as information frictions fall.

A genuinely non-obvious reconciliation of two literatures:

- AI as information-processing technology (Cheng et al. 2025; Bertomeu et al. 2025).
- Disagreement as heterogeneous attention / interpretation (Cookson et al. 2023; Acemoglu et al. 2025).

# Comment 1 — "Tone" may not be "belief"

The whole paper hinges on FinBERT-tone-of-prompt $\approx$ investor belief. Three reasons to worry:

(a) **Task framing.** "Could receivables create earnings-quality problems?" produces negative tone without expressing a view. So does "summarize the bear case."

(b) **Prior framing, not posterior.** $|\text{tone}|$ rising 25% may mean prompts get more pointed as the conversation narrows, not that beliefs become more extreme.

(c) **Strategic anti-belief prompting.** Sycophancy is now common knowledge. Sophisticated users may deliberately prompt *against* their belief to elicit a critical answer $\Rightarrow$ tone may be *systematically inversely* related to belief.

**Suggestions.**

- Classify a sample of prompts into belief / info request / risk inquiry / stance request.
- Show the main results survive on the "belief statement" subset.
- Separately analyze prompts that explicitly request a stance.

# Comment 2 — Is between-topic dispersion really "disagreement"?

The headline mechanism: AI raises *between-topic* dispersion.

But if I care about ESG risk and you care about Q3 earnings, we are not disagreeing about firm value — we are **computing different objects**.

The Harris–Raviv / Kandel–Pearson link from disagreement to trading volume assumes disagreement about the *same* fundamental. Cross-topic variance $\neq$ disagreement-driven trade.

**Suggestions.**

- Frame the result as **heterogeneous attention allocation**, not "disagreement."
- Map topic-specific tones to a common value scale (e.g., have AI extract a directional valuation signal from each response) and decompose dispersion on that scale.
- Be explicit about when between-topic variance translates into trading-relevant disagreement.

# Comment 3 — Topic stickiness: behavioral, or something else?

The paper attributes topic persistence to behavioral confirmation seeking. Three alternatives deserve weight:

(a) **Model architecture.** LLMs maintain conversational context. Continuing the same topic is the path of least resistance *for the model*, not necessarily a user bias.

(b) **Rational task focus.** A user investigating governance risk *should* stay on governance risk. Value of additional rounds may show up in risk calibration, not 1–5 day return predictability. Absence of incremental return predictability $\neq$ absence of incremental value.

(c) **Platform specialization** (see Comment 4).

**Suggestions.**

- Use a softer label: "persistent topic attention."
- Test stickiness on subjective topics (ESG, sentiment) vs. factual ones (balance sheet).
- If observable, compare reset-context conversations to maintained-context ones.

# Comment 4 — One platform, no trades

The paper observes **one** Chinese GenAI platform; no trades observed.

In practice retail investors multi-platform: DeepSeek, Doubao, Kimi, Wenxin, ChatGPT-via-VPN — each with different strengths.

**Consequences.**

- **Topic stickiness may be an artifact of platform specialization.** If a user goes to Platform A for fundamentals and Platform B for ESG, queries on Platform A look highly sticky — not because the user is sticky.
- **Belief updating** attributed to "this AI" may come from another tool used between rounds.
- **No trades** $\Rightarrow$ the chain tone $\rightarrow$ belief $\rightarrow$ action is asserted, not closed.

**Suggestions.**

- Acknowledge the bias direction explicitly.
- Test whether stickiness varies with sophistication proxies — the direction of bias is not obvious (sophisticated users may multi-platform more, or simply ask more topics), but the test itself would be informative.

# Comment 5 — The user-updating regression

Table 6 regresses

$$\Delta\text{Tone}^{\text{Human}} = \text{Tone}_2^{\text{Human}} - \text{Tone}_1^{\text{Human}} \quad \text{on} \quad \text{Gap} = \text{Tone}_1^{\text{AI}} - \text{Tone}_1^{\text{Human}}.$$

**(a) Mechanical mean reversion.** $\text{Tone}_1^{\text{Human}}$ enters with a $-$ sign on *both* sides. Extreme initial tones revert mechanically $\Rightarrow$ positive coefficient even with no AI influence.

**(b) Strategic-prompting reinterpretation.** A user who asks for the bear case and receives a bearish answer has not been persuaded — she got what she engineered for. The "alignment" may partly reflect a continuing stance, not belief updating.

**Suggestion.**

- Regress $\text{Tone}_2^{\text{Human}}$ on $\text{Tone}_1^{\text{AI}}$ and $\text{Tone}_1^{\text{Human}}$ separately — test, don't impose, the restriction.

# Comments 6 & 7 — Two shorter points

**Comment 6 — The U.S. ChatGPT test is suggestive, not clean.**

- PostGPT dummy (Dec 2022+) coincides with rates, macro, AI hype trade, analyst-behavior shifts.
- Closer to a pre/post test than a true diff-in-diff — no cross-sectional treatment intensity.
- Strengthen with: disclosure complexity, low-coverage firms, high-retail firms, AI-related search interest.

**Comment 7 — "Noise" is not the same as hallucination.**

- Declining within-topic AI tone dispersion shows *consistency*, not *accuracy*.
- A model can be consistently wrong, or give similar-tone answers with different factual errors.
- Strengthen with: manual factual audit, comparison to filings, retrieval vs. non-retrieval responses.

# Minor comments — two cross-sectional tests worth running

**A. By prompt sophistication.**

- Skilled prompters likely elicit different responses than novices, and interact differently (multi-platform; strategic stance-engineering — see Comments 1c and 4).
- Is dispersion higher among less-sophisticated users?
- Could the between-topic effect be largely a *novice*-prompt phenomenon?
- Proxies: prompt length, structure, lexical complexity, account age.

**B. By firm size / analyst coverage.**

- Theory: AI should matter most for small, under-covered firms (high retail processing costs, sparse coverage).
- Does the dispersion increase concentrate in small-cap, low-coverage firms?
- Also the natural cross-section for tightening the U.S. design (Comment 6).

# Wrap-up

- A timely, important, and interesting paper.
- Main asks: validate tone-as-belief; reframe between-topic dispersion; soften the behavioral interpretation of stickiness; tighten identification on belief updating and the U.S. test.

**Thank you.**