

Do AI Trading Agents Herd? Experimental Evidence and Financial Stability Implications

AUTHORS: Anne Artrip and Seung Jung Lee

SUMMARY: This paper investigates whether AI-powered trading agents exhibit herd behavior, a phenomenon where traders follow the crowd rather than relying on their own information. The authors use laboratory-style experiments with large language models to replicate classic studies on herd behavior in trading decisions. Their findings suggest that AI agents rely more on private information than market trends, unlike human traders. However, AI can be guided to herd optimally when explicitly instructed to maximize profits.

METHOD: The study employs experimental methods using large language models in a controlled environment, replicating classic studies on herd behavior. The sample includes both AI agents and human traders. Data are collected through simulated trading scenarios to observe behavior under different conditions.

KEY FINDINGS:

- AI agents rely less on market trends and more on private information compared to human traders.
- AI can be induced to herd optimally when explicitly guided to make profit-maximizing decisions.
- AI agents inherit some elements of human conditioning and bias.
- The results highlight the potential for AI to amplify or mitigate financial instability.

IMPLICATION FOR TRADING: The research suggests that while AI can reduce herd behavior by relying more on private information, it can also be manipulated to herd when profit maximization is the goal. This has significant implications for financial stability, as AI's behavior can either stabilize or exacerbate market volatility. Traders and regulators should be aware of these dynamics to better understand and manage the risks associated with AI in financial markets.

Do AI Trading Agents Herd? Experimental Evidence and Financial Stability Implications *

Anne Artrip^{a,b} and Seung Jung Lee^a

^a*Board of Governors of the Federal Reserve System*

^b*Federal Reserve Bank of Richmond*

First draft: September 25, 2025.

This draft: September 18, 2026 .

This paper provides micro-foundations for understanding how AI-powered trading affects market dynamics. We conduct laboratory-style experiments using large language models to replicate classic studies on herd behavior in trading decisions. Our results show that AI agents rely less on market trends and more on private information compared with human traders. However, exploring variations in the experimental settings reveals that AI agents can be induced to herd optimally when explicitly guided to make profit-maximizing decisions. In other experimental variations, we show that AI agents are not purely algorithmic, but have inherited some elements of human conditioning and bias.

Keywords: Herd behavior, large language models, AI-powered traders, financial markets, financial stability.

JEL Codes: C90, D82, G11, G14, G40.

* The authors thank Jeffrey S. Allen, Marco Cipriani, Ricardo Correa, Eddie Gerba, Glenn Harrison, Erik Heitfield, Sophia Kazinnik, Dan Li, Molly Mahar, Nitish Sinha; discussants Haoyang Liu, Summer Liu, Tara Sinclair, and Todd Swarthout; and participants at the Financial Stability Workshop at the Federal Reserve Board, the 2025 “AI and the Economy” at the Federal Reserve Bank of San Francisco, 19th International Behavioural Finance Conference, the 2025 “Improving Financial Regulation for an Evolving Environment” workshop, the 4th CEMLA/Dallas Fed Financial Stability Workshop, the ECONDAT knowledge sharing session, the 2026 Annual IJCB Research Conference, and the XVI Meeting of Heads of Financial Stability participants for valuable feedback. The views expressed in this paper are solely those of the authors and do not reflect the opinions of the Federal Reserve Bank of Richmond or the Board of Governors of the Federal Reserve System. All errors are the authors’. Contact information: anne.artrip@rich.frb.org (corresponding author), seung.j.lee@frb.gov. Anne Artrip was previously known as Anne Lundgaard Hansen and this paper was previously circulated under the title “*Financial Stability Implications of Generative AI: Taming the Animal Spirits.*”

...[T]here is the instability due to the characteristic of human nature that a large proportion of our positive activities depend on spontaneous optimism rather than mathematical expectations [...]. Most, probably, of our decisions [...] can only be taken as the result of animal spirits—a spontaneous urge to action rather than inaction, and not as the outcome of a weighted average of quantitative benefits multiplied by quantitative probabilities.

— John Maynard Keynes

1. Introduction

Human irrationality is a key driver of the build-up of financial vulnerabilities, contributing to asset price bubbles and banking crises. History offers numerous examples, including Tulip Mania in the 17th century, the South Sea Bubble, the dot-com boom, the 2008 financial crisis, the 2010 Flash Crash, and the GameStop short squeeze. A well-established body of research highlights the role of psychological and emotional factors, coined *animal spirits* or *irrational exuberance*, in these periods of boom and bust (Angeletos et al., 2018; Grauwe, 2012; Shiller, 2005).

Understanding the role of animal spirits for financial stability is already a challenge, given the unpredictable nature of human behavior. Recently, a new and unknown agent has entered the equation: decisions powered by generative AI. Humans increasingly rely on AI for information gathering and decision making, whether as a co-pilot or as autonomous agents.¹ As generative AI is reshaping workflows across institutions and individuals, the question arises: How might the increased reliance on generative AI impact financial stability? Specifically, will generative AI exaggerate or dampen the role of animal spirits in the build-up of financial vulnerabilities?

Two competing hypotheses emerge. On the one hand, AI is fundamentally algorithmic, operating within a set of instructions and grounded in logic and rational decision making.² If AI-guided decisions replace human intuition, the result could be a reduction in the influence of animal spirits, leading to

1 Hartley et al. (2024) conducted a survey covering the U.S. labor market showing that the percentage of workers adopting LLMs steadily increased from 30% to 45% between end of 2024 and mid 2025. These adoption rates are massive in the context of an earlier survey from 2018 showing that less than 6% of firms used any AI-related technology, i.e., automated-guided vehicles, machine learning, machine vision, natural language processing, or voice recognition (McElheran et al., 2023). More specifically, based on data of ChatGPT outages, Cheng et al. (2025) show that ChatGPT is used by investors in ways and to an extent that impact trading decisions and market outcomes.

2 This algorithmic foundation suggests that AI systems may enhance human judgment by providing consistent assessments less prone to biases. The seminal paper by Kleinberg et al. (2018) argues that machine-learning algorithms can improve human decision making in the context of bail decisions, especially if carefully integrated into an economic framework. Similarly, Li et al. (2024) demonstrates how AI-enabled credit scoring models can increase loan approval rates for under-served populations while simultaneously reducing default rates, primarily through the algorithmic processing of weak signals that humans might overlook or inconsistently evaluate.

more stable financial markets. On the other hand, generative AI models, such as large language models (LLMs), are trained on vast amounts of data, sourced from both rigorous materials, such as academic research, and the, at times, chaotic discourse of social media platforms such as Twitter (X) and Reddit. Consequently, generative AI may inherit and even amplify human biases and irrational tendencies (Hayes et al., 2024; Jiang et al., 2023; Koralus & Wang-Maścianica, 2023; Zhu & Griffiths, 2024). Moreover, many AI models undergo reinforcement learning from human feedback (see Wang et al., 2024 for a survey), optimizing for engagement and persuasion rather than pure rationality. This suggests that instead of mitigating the build-up of financial vulnerabilities, AI could exacerbate financial turbulence driven by animal spirits. Finally, Danielsson and Uthemann (2024) argue that AI adoption will likely cause more intense future crises due to AI’s ability to respond quickly to shocks. The net effect of AI’s involvement in financial decision making is therefore unclear.

This paper explores these competing perspectives, examining the implications of the expanding role of AI in economic decision making for financial stability. We focus on the potential financial vulnerabilities driven by herding behavior, where investors ignore private signals and mimic the decisions of others. The risk of increased herding is among the top cited concerns associated with the adoption of generative AI according to the International Monetary Fund (2024) Global Financial Stability Report. Herd behavior is a well-documented cause of asset price bubbles (Galariotis et al., 2016; Hsieh et al., 2020). Even in cases where herding is optimal, this type of behavior can contribute to financial stability events by increasing market volatility and accelerating price movements which can reverse quickly in the event of new information arriving (Bikhchandani & Sharma, 2001). Optimal herding can also drive prices away from fundamentals because signals or interpretation of information can be wrong (Chamley, 2003).

We conduct laboratory-style experiments using LLMs to replicate classic studies on herd behavior in trading decisions. These experiments provide novel insights into the behavioral tendencies of AI agents, laying a micro-foundation for future work on financial stability in an AI-powered economy. Just as laboratory experiments with human subjects established behavioral regularities that inform theoretical models and field studies, our AI laboratory provides behavioral benchmarks. Our micro-level approach is motivated by Horton (2023), who argues that LLMs can be treated as agents whose decisions and behavior can be studied in parallel to studies of human behavior. We stress that we are not attempting to model a realistic financial market, but rather to zoom in on the behavior of LLMs in a controlled setting.³ In particular, our setting allows us to observe detailed data on each decision made by AI agents, including

3 While a complementary study of financial stability implications of generative AI for macro-level outcomes is certainly of interest, we conjecture that such an undertaking would be difficult to achieve. First, there are the usual problems of disentangling effects from human behavior from confounding factors and distinguishing intentional herding from situations where the private information of many traders happen to coincide, coined as “spurious herding” in Bikhchandani and Sharma (2001). Second, the world is still in early stages of generative AI adoption with impacts yet to be seen.

their reasoning.

We focus on the experiment by Cipriani and Guarino (2009), which investigates herd behavior at the extensive margin, i.e., regarding the binary decision to participate in trading or not rather than on how much to trade conditional on participation. This distinction is crucial for financial stability, as extensive margin coordination creates discontinuous regime shifts that can trigger systemic crises. The Cipriani and Guarino (2009) experiment is grounded in an established model of herd behavior (Avery & Zemsky, 1998), which has been rigorously tested in laboratory-style experiments (Cipriani & Guarino, 2005; Drehmann et al., 2005). Whereas the previous experimental literature conducted experiments with undergraduate students, Cipriani and Guarino (2009) recruited financial market professionals. These details make their study particularly interesting as a human benchmark for our purpose. After all, it is the actions of financial market professionals—not students—that shape real-world market dynamics and impact the stability of the financial system.

We replicate the Cipriani and Guarino (2009) experiments using trading decisions of LLMs (which we refer to as *AI agents*) in place of decisions made by human participants. Our implementation closely following the original experimental design. Specifically, we prompt the LLMs with instructions that mirror those given to financial professionals in the human study. This setup allows us to compare LLM decision making with the human results taken directly from Cipriani and Guarino (2009): the “AI laboratory” versus the “human laboratory.” To generalize our results as much as possible, we use four different LLMs⁴ and average the results across models.

Our results show that, in contrast to human participants, AI agents mainly rely on their private information and less on market trends in their trading decisions. Across different parameterizations of the experiment, AI agents made private-information-driven decisions between 61-97% of the time, substantially exceeding the 46-51% range observed in human participants. The AI laboratory also exhibited fewer information cascades, where investors trade based on the actions of other investors rather than relying on their own private information. Specifically, information cascades occurred between 0-9% of the AI decisions, compared to around 20% for humans. Notably, when AI agents did engage in cascade trading behavior, they traded against market trends (contrarian behavior) rather than following the herd. In addition, we show that AI agents do not exhibit erroneous behavior, which contrasts with the results of the experiments conducted with human participants. We interpret these results as early indications that a future where investors are more impacted by advice generated by LLMs can potentially involve fewer asset price bubbles arising from herd behavior.

However, studying the rationales provided by LLMs alongside the trading decisions reveals that AI

⁴ These LLMs are: Anthropic’s Claude 3.5 Sonnet and Claude 3.7 Sonnet models, Meta’s Llama 3 Instruct 70B model, and Amazon’s Nova Pro model.

agents rely too heavily on their private information. As market trends can reflect the private information of others, it *can* be optimal from a profit-maximizing perspective to take trading history into account when making trading decisions. The literature therefore distinguishes between optimal and suboptimal herding: Optimal herding represents a fundamental mechanism through which individually optimal decisions can generate collective fragility in financial markets. In the canonical models of Bikhchandani et al. (1992) and Banerjee (1992), optimal herding occurs when mimicking others represents an optimal response to superior information possessed by early movers. Once a sufficient number of agents have acted in one direction, subsequent investors optimally ignore their own private signals, leading to collective behavior that may diverge from fundamental values. Suboptimal herding occurs when investors follow the crowd even when this leads to lower expected profits than decisions relying on private information. This behavior stems from cognitive biases, reputational concerns (Scharfstein & Stein, 1990), and misaligned incentives. One recent example of suboptimal herding behavior is the “meme stock mania” of 2021, exemplified by GameStop and AMC, where retail investors collectively moved markets by following social trends rather than fundamental valuations.

We show that AI agents fail to acknowledge opportunities for optimal herding, leading to occasional suboptimal choices. We therefore implement AI agents that are prompted with additional guidance on optimal decision making; instructions that their human counterparts did not receive in the original field experiment. This experiment serves as a proxy for the fine-tuning that financial institutions would likely implement if adopting generative AI for powering trading. We show that these optimal AI agents do engage in cascade trading when optimal, but that they still remain reluctant to herding.

Finally, we explore variations of the experiment to examine whether different conditions lead to stronger evidence of herd behavior. Unlike the original study, scaling up the length of the experiments and modifying their parameters is both cost-effective and efficient with LLMs. We generally find that our results are robust to these variations. As an exception, we test the impact of re-labeling the private signals that participants receive during the experiment. The original experiment uses neutral color-coding for these signals. We test outcomes using non-neutral colors, such as green and red. Using “green” to indicate a high probability of a high asset value and “red” to indicate a high probability of a low asset value yield similar results as the baseline experiment. However, reverting the labeling such that “red” (“green”) signals a high probability of a high (low) asset value, which is counterintuitive given human conditioning, the LLMs generate similar trading decisions as the human traders. AI agents are therefore not algorithmically rational, following a well-defined set of rules, but have inherited some elements of human intuition and bias. This finding is consistent with a growing literature showing that LLMs can replicate human errors and biases (Argyle et al., 2023; Bybee, 2023; Hansen et al., 2025).

Our results suggest that AI agents exhibit less herd behavior than human financial professionals, a

finding with significant implications for future financial stability as generative AI gains traction in market decision making. Specifically, the reduced tendency to herd could potentially lead to less extreme market movements and fewer asset price bubbles, contributing to greater overall financial market stability. However, we acknowledge that the introduction of AI agents could fundamentally alter market dynamics in ways that are not yet fully understood. Continued research and adaptive regulatory approaches to maintain financial stability in an AI-augmented financial landscape is therefore warranted.

We proceed as follows. Section 2 reviews the literature. Section 3 outlines the theoretical model that underpins the experimental design. Section 4 describes the human laboratory in which the experiment was conducted in Cipriani and Guarino (2009), and how we adopt this setting with LLMs. Section 5 presents the main results. Section 6 discusses implications of our results for financial stability. Conclusions follow in Section 7.

2. Literature

Previous work has emphasized that AI-powered trading could lead to more correlated behaviors, e.g., stemming from collusion (Dou et al., 2025),⁵ market concentration (Bank of England Financial Policy Committee, 2025), and model monoculture (Danielsson & Uthemann, 2024; Financial Stability Board, 2024). While these concerns are certainly relevant to consider, they are not the focus of our work. We aim to provide insights into behavioral aspects of LLMs and their tendency to herd in trading decisions, and leave the impact of AI adoption on the dynamics among market participants for future research.

Our work contributes to a growing literature on the behavior of LLMs. While our study focuses on herd behavior and financial stability, other works have examined other types of behavior and departures from rationality. Most relevant to our approach are Henning et al. (2025) and Muntoni (2025), who conduct asset pricing experiments with LLM traders. Henning et al. (2025) demonstrate that AI agents tend to price assets near fundamental values. As in our work, they conclude that AI adoption has the potential to enhance financial stability by dampening the likelihood of asset price bubbles. Their test is, however, fundamentally different from ours. In Henning et al. (2025), agents choose whether to invest cash in a risky asset under full information about expected dividends and interest rates, which directly facilitates computation of the fundamental value. This setup seeks to test rationality in the context of price setting rather than behavioral aspects of asset price bubbles, which is our focus. Muntoni (2025) focuses on the formation of expected future prices given past realized prices. Consistent with our results, he finds evidence that AI agents generate prices that are less volatile and closer to fundamental

⁵ Collusion involves strategic and deliberate coordination. In contrast, herding is uncoordinated imitation arising from informational spillovers or psychological biases.

values than human traders. However, in contrast with our findings, these conclusions are not robust to different model configurations such as temperatures and prompt instructions. Anand et al. (2025) also studies investment behavior of LLMs, using a mutual fund redemption problem. Similar to our conclusion, they find that LLM traders weaken coordination of market movements. Their identified channel is, however, different from ours: Whereas our experiment show that AI traders primarily follow private information, which weakens coordination, the mutual fund redemption experiment show that AI traders exhibit heterogeneous beliefs. Cartea et al. (2026) show that AI agents participate in both rational and irrational speculative bubbles. Supporting our findings, they show that bubble behavior arises mainly due to speculation, not due to herding behavior, and that AI agents engage in irrational bubbles due to simplified views of market dynamics.

Chen et al. (2023) study the economic rationality of GPT models by conducting revealed preference experiments, where models are prompted to make decisions under budget constraints. Similar to our results, although in a different aspect of the term rationality, the authors conclude that AI agents tend to exhibit more rational behavior than humans. Liu et al. (2025) confirm that LLMs are more rational than humans using a large data set of human decisions in risky choice problems. Bini et al. (2025) consider LLMs in a wide set of behavioral experiments covering both preferences and belief generation. In this context, our setting most closely resemble a belief-related experiment, as it involves the computation of the expected value of the asset. Consistent with our work, Bini et al. (2025) show that advanced LLMs often generate rational responses. del Rio-Chanona et al. (2025) focus on laboratory experiments related to price expectations and deviations from rational expectations. They emphasize the importance of the interactions of different AI agents and retaining memory across time periods; both elements that we include into our AI laboratory setting as well. While they conclude that LLMs are not strictly rational in their expectation formation, they find that LLMs generate less variability in their responses compared with humans. Similar patterns are observed in our results.

While these studies, like ours, mainly emphasize differences between the behaviors of humans and AI agents, some papers emphasize their similarities and argue that LLMs can be used to simulate human outcomes. For example, Horton (2023) argues that LLMs can give human-like responses and suggests that they can be used conduct pilot experiments to calibrate experimental designs before testing on human beings. Hansen et al. (2025) and Jha et al. (2024), and Zarifhonarvar (2024) show that LLMs can be used to simulate economic surveys. The literature emphasizes that LLMs in some contexts exhibit human biases such as risk aversion and loss aversion (Jia et al., 2024; Ross et al., 2024). Along the same lines, Hua et al. (2024) show that LLMs often deviate from rational decisions in game theoretic experiments. Characterizing the exact distinction between human and AI decision making remains an open question.

Finally, our work is particularly important as LLMs start to play an increasingly larger role in financial

market decisions. The literature showcases the application of LLMs in investing. For example, Lopez-Lira (2025) simulates a stock market using LLMs and argues that this framework can be used to conduct counterfactual experiments. Lee et al. (2025) argue that LLMs suffer from confirmation bias in the realm of investment. Specifically, the authors show that LLMs exhibit a preference for large-cap stocks and contrarian trading strategies. And Fedyk et al. (2024) survey the investment preferences of human and AI agents, finding that AI demonstrates demographic biases that can be overcome with demographically-seeded prompts.

3. Experimental design

This section presents the model and theoretical predictions outlined in Cipriani and Guarino (2009). The model is based on Avery and Zemsky (1998).

3.1. Theoretical model

The model describes a financial market with one risky asset and discrete trading periods indexed by $t = \{1, 2, \dots\}$. Before the first trading period, there is a ρ probability of an *information event*, which changes the fundamental value of the asset in either direction. In each trading period, some traders are informed and receive a private signal on the value change, while others do not. The model characterizes different types of trading behaviors based on whether informed traders act according to their private signal (coined *rational* behavior⁶) or ignore their signal (*cascade* behavior).

The asset’s fundamental value belongs to the discrete set $v \in \{0, 50, 100\}$. Specifically, if there is no information event (with probability $1 - \rho$), the value is equal to its unconditional expected value given equal probabilities, i.e., $v = 50$. An information event (occurring with probability ρ) pushes the value to zero or 100, with the following probability distribution: $\Pr(v = 0) = \Pr(v = 100) = 0.5$. The asset trades at a price p , which is set by the market maker according to Bayesian updating as we detail below.

Traders act sequentially with only one trader randomly chosen to trade in each trading period.⁷ In each period t , the chosen trader executes an action x_t , which is to buy one unit of the asset ($x_t = \text{buy}$), sell one unit of the asset ($x_t = \text{sell}$), or not trade ($x_t = \text{no trade}$). Since the experiment focuses on the trading decision of one unit of the asset rather than the decision of the number of units to trade, our focus is on herding at the extensive margin. If there is no information event, all traders are uninformed *noise traders*,

6 Note that “rational” behavior does not imply “optimal” behavior. This distinction is discussed further in Section 3.2. We adopt the term “rational” to describe actions purely driven by private information for consistency with the herding literature.

7 This structure simulates the mechanics of central limit order book trading.

who trade based on exogenous probabilities, i.e., $\Pr(x_t = \text{sell}) = \Pr(x_t = \text{buy}) = \Pr(x_t = \text{no trade}) = 1/3$. In the case of an information event, the chosen trader is *informed* with probability μ (and a noise trader with probability $1 - \mu$). An informed trader receives a signal $s_t \in \{\text{white}, \text{blue}\}$, which is tied to the asset value in the following way:

$$\Pr(s_t = \text{white} \mid v = 100) = \Pr(s_t = \text{blue} \mid v = 0) = 0.7. \quad (1)$$

That is, a white signal can be interpreted as a *good* signal, indicating that the information event resulted in a high asset value, whereas a blue signal is *bad* in the sense that it increases the probability of a zero asset value. In addition to the signal s_t , an informed trader also observes the trading history h_t , and therefore forms beliefs about the asset value based on the conditional expected value given s_t and h_t : $\mathbb{E}(v \mid s_t, h_t)$. The realized payoff is equal to $v - p$ if the trader chooses to buy the asset, $p - v$ if the trader chooses to sell, and zero if the trader chooses not to trade. We assume that the informed trader is risk-neutral and seeks to maximize expected payoff given s_t and h_t .

A market maker facilitates exchanges with the traders and sets the price of the asset given the history of trades for periods up to $t - 1$, $h_t = \{x_1, x_2, \dots, x_{t-1}\}$ for $t > 1$. $h_1 = \emptyset$. Specifically, the price is determined as the expected asset value given h_t : $p_t = \mathbb{E}(v \mid h_t)$.⁸ In the first trading period, with no trading history, the price is equal to its unconditional expected value: $p_1 = \frac{1}{2}100 = 50$. At $t > 1$, the price is given by the expected asset value conditional on the history of trades:

$$p_t = 100 \Pr(v = 100 \mid h_t) + 0 \Pr(v = 0 \mid h_t) = 100q_t \quad (2)$$

where $q_t = \Pr(v = 100 \mid h_t)$ is determined using Bayesian updating:⁹

$$q_t = \Pr(v = 100 \mid x_{t-1}, h_{t-1}) \quad (3)$$

$$\begin{aligned} &= \mathbb{I}_{(x_{t-1}=\text{buy})} \left[\frac{(0.7\rho\mu + (1 - \rho\mu)\frac{1}{3}) q_{t-1}}{(0.7\rho\mu + (1 - \rho\mu)\frac{1}{3}) q_{t-1} + (0.3\rho\mu + (1 - \rho\mu)\frac{1}{3}) (1 - q_{t-1})} \right] + \\ &\quad \mathbb{I}_{(x_{t-1}=\text{sell})} \left[\frac{(0.3\rho\mu + (1 - \rho\mu)\frac{1}{3}) q_{t-1}}{(0.3\rho\mu + (1 - \rho\mu)\frac{1}{3}) q_{t-1} + (0.7\rho\mu + (1 - \rho\mu)\frac{1}{3}) (1 - q_{t-1})} \right] + \\ &\quad \mathbb{I}_{(x_{t-1}=\text{no trade})} q_{t-1}. \end{aligned} \quad (4)$$

8 There is only one asset price, i.e., the model assumes a zero bid-ask spread. This assumption was imposed by Cipriani and Guarino (2009) to simplify the laboratory experiment.

9 The term $(1 - \rho\mu)\frac{1}{3}$ represents the probability a buy or sell comes from a noise trader, who buys, sells, and chooses not to trade with equal probability. The term $\rho\mu$ is the probability that a trader is informed, given by the probability that an information event occurred (ρ) times the probability that a trader is informed given an informed event (μ).

3.2. Theoretical predictions

This section presents the theoretical predictions for how informed traders act according to the model. Informed traders make decisions by comparing the price of the asset to the expected value given the signal and trading history:

$$x_t = \begin{cases} \text{buy} & \text{if } p_t < \mathbb{E}(v|s_t, h_t) \\ \text{sell} & \text{if } p_t > \mathbb{E}(v|s_t, h_t) \\ \text{indifferent} & \text{if } p_t = \mathbb{E}(v|s_t, h_t) \end{cases} \quad (5)$$

When indifferent, traders may buy, sell, or not trade; their payoff will be the same regardless of their action. Their expected value is given for each signal as follows:

$$\mathbb{E}(v|s_t = \text{white}, h_t) = 100 \left[\frac{0.7q_t^*}{0.7q_t^* + 0.3(1 - q_t^*)} \right], \quad (6)$$

$$\mathbb{E}(v|s_t = \text{blue}, h_t) = 100 \left[\frac{0.3q_t^*}{0.3q_t^* + 0.7(1 - q_t^*)} \right] \quad (7)$$

where $q_t^* = \Pr(v = 100|x_{t-1}, h_{t-1}, \rho = 1)$, which can be computed from (4) above. For the informed trader, the relevant probability of the high asset value conditional on the trading history sets $\rho = 1$ because the informed trader, by definition, knows with certainty that an information event occurred. The discrepancy between q_t , the probability of a high asset value from the perspective of the market maker, and q_t^* , the corresponding probability from the perspective of an informed trader, can lead to optimal information cascades.

The model characterizes different types of behavior of informed traders. We label these behaviors consistently with the herding literature (see, e.g., Cipriani and Guarino, 2005 and Cipriani and Guarino, 2009), i.e., as follows:

Rational: The informed trader chooses to buy upon receiving a white (*good*) signal and sell upon received a blue (*bad*) signal.

Partial rational: The informed trader follows rational behavior upon receiving one signal and to not trade upon receiving the other signal, e.g., buy upon receiving a white (*good*) signal and no trade upon received a blue (*bad*) signal.

Cascade trading: The informed trader chooses the same trading action (buy or sell) regardless of the private signal. If the trader chooses to buy (sell) when the trading history is dominated by buy-actions (sell-actions), i.e., act following the majority action of previous traders, the trader engages

in *herd behavior*. If the trader chooses to buy (sell) when the trading history is dominated by sell-actions (buy-actions), i.e., acting against the majority of previous traders, the trader engages in *contrarian behavior*.

Cascade no trading: The informed trader chooses not to trade regardless of the private signal.

Error: The informed trader chooses to buy upon receiving a blue (*bad*) signal and sell upon receiving a white (*good*) signal.

The last type of behavior is always sub-optimal and is interpreted as an error if observed. However, it can be optimal for traders to engage in cascade behavior, depending on the parameterizations of the model. The laboratory experiments in Cipriani and Guarino (2009) follow two different parameterizations of the model, referred to as treatments.

In the first treatment (Treatment I), there is no uncertainty about whether an information event occurs, i.e., $\rho = 1$. In addition, all traders are informed, i.e., $\mu = 1$. Hence, $q_t = q_t^*$, and it follows that:

$$\mathbb{E}(v|s_t = \text{white}, h_t) = 100 \frac{0.7q_t}{0.7q_t + 0.3(1 - q_t)} > 100q_t = p_t$$

and

$$\mathbb{E}(v|s_t = \text{blue}, h_t) = 100 \frac{0.3q_t}{0.3q_t + 0.7(1 - q_t)} < 100q_t = p_t.$$

Hence, regardless of the history of trades, a trader's expected value given their private signal is always on the same side of the market price as their signal. Therefore, it is always optimal for traders to follow their private signals. As a result, each trade reveals new information, continuously updating the market price. This prevents the formation of information cascades, as traders never have an incentive to ignore their private information in favor of following the actions of others.

In the second treatment (Treatment II), there is uncertainty both about whether an information event occurs and the proportion of informed traders. Cipriani and Guarino (2009) set $\rho = 0.15$ and $\mu = 0.95$, i.e., an information event occurs with 15% probability and the probability that a trader receives a private signal on the information event is a slightly smaller than one.

With event uncertainty, it can be optimal for traders to engage in cascade behavior. The reason is that there is information asymmetry between informed traders and the market maker. Upon receiving a private signal, the informed trader knows with certainty that an information event has occurred and that the history of trades comes from an informed trader with probability $\mu = 0.95$. In contrast, not knowing whether an information event has occurred, the market maker believes that the traders are informed with probability $\rho\mu = 0.15 \cdot 0.95 = 0.14$. This asymmetry leads the market maker to update the asset price

more conservatively than informed traders update their beliefs. After a sequence of buy orders, the gap between traders' expectations and the market price can widen. Eventually, a trader's expectation may exceed the market price even with a contradictory signal: $\mathbb{E}(v|s_t = \text{white}, h_t) > \mathbb{E}(v|s_t = \text{blue}, h_t) > p_t$. At this point, the trader will ignore their private information and follow the herd.¹⁰ However, because the market maker updates his expectation by less than the informed traders, it will never be the case that, after a history of buys, the expectation of a trader will be below the price for both signal realizations, i.e., $p_t > \mathbb{E}(v|s_t = \text{white}, h_t) > \mathbb{E}(v|s_t = \text{blue}, h_t)$. As a result, an informed trader will never engage in contrarian behavior. Analogous arguments apply to a sequence of sell orders.

At the extreme, the market maker does not update the price at all such that the price remains at the unconditional expected value throughout all trading periods. Cipriani and Guarino (2005) conducted an experiment with this setting (without event uncertainty) among undergraduate students. We shall refer to this setting as Treatment III. In this parametrization, optimal herding arises when there is a trade imbalance greater than or equal to two (Bikhchandani et al., 1992); see Cipriani and Guarino (2005) for intuition. Since this experiment was not conducted among financial market professionals, we shall focus less on this parametrization in our results.

Optimal behavior: To summarize, the model predicts the following behavior in the two treatments:

Treatment I (price updating; no event uncertainty): Traders always trade according to their private signal, preventing the formation of cascades.

Treatment II (price updating; event uncertainty): An information cascade occurs with positive probability. Herding is optimal when prices are below the expected value conditional on both signals, but it is never optimal to engage in contrarian behavior.

Treatment III (no price updating; no event uncertainty): Herding is optimal after a trade imbalance higher than or equal to two.

4. Laboratory setup

Cipriani and Guarino (2009) implemented the experiment among financial market professionals. We adopt their *human laboratory* setting as closely as possible, replacing human participants with AI agents.

¹⁰ Optimal herding behavior is temporary. When traders herd, the private signals are not reflected in the prices. However, the market maker continues to update beliefs about whether an information event has occurred, causing prices to keep moving, albeit slowly. Eventually, the price may move enough to make private information relevant again, breaking the herd behavior.

Then we compare our results from this *AI laboratory* with the human results from Cipriani and Guarino (2009). This section describes the human and AI laboratories.

4.1. Human laboratory

The human experiment was conducted with 32 participants working for financial institutions in London. The participants were divided into four groups of eight; each group formed one session.

In each of the four sessions, the experiment was repeated for two practice rounds followed by first eight rounds implemented with the parametrization in Treatment I and then eight rounds with the Treatment II parametrization. Before each treatment, participants were given written instructions. They were informed that everyone received the same set of instructions, and were given the opportunity to ask clarifying questions which were answered privately. The timelines for each session and each round are included in Electronic Appendix. We refer to Cipriani and Guarino (2009) for further details.

4.2. AI laboratories

We adopt the human experiment in our AI laboratory, where human participants are replaced by AI agents. To model AI agents, we use a suite of LLMs and apply model averaging to get an all-compassing view of the behavioral patterns of AI-powered trading. Specifically, we use Anthropic’s Claude 3.5 Sonnet and Claude 3.7 Sonnet models (hereafter Claude 3.5 and 3.7), Meta’s Llama 3 Instruct 70B parameter model, and Amazon’s Nova Pro model. Importantly, we are not claiming that the errors across LLMs are zero in expectation and that model averaging therefore achieves unbiased estimates. Rather, model averaging serves to represent the views of multiple LLMs in a single measure.¹¹ We mainly implement the models with a moderate temperature of 0.7, balancing creativity with determinism.¹² Robustness checks confirm that the choice of temperature does not impact the conclusions of our experiments. The Claude 3.7 model is implemented in *extended thinking mode*, requiring a temperature of 1.0. This setting activates extended reasoning capability, where the model iterates in multiple steps to reach the “best” perceived answer.

We follow the setup of the human laboratory described above as closely as possible. For example, similar to human participants, we presented an LLM (Claude 3.5) with written instructions and gave the model the opportunity to ask clarifying questions. We used this model feedback to improve the instructions. However, some adjustments are necessary to accommodate differences between human

11 In cases where conclusions vary across LLMs, we do report results for each model separately.

12 The temperature adjusts how the model weighs its prediction for the next token. A lower temperature makes the model focus more on its top choices, while a higher temperature gives it more freedom to consider less likely options, affecting how predictable or creative the output becomes.

and AI agents. First, practice rounds are redundant, and we completely separate the two treatments to avoid confusing the models. Second, we explicitly provide memory to the AI agents in each trading period, by listing the executed trades along with the history of actions and reasoning for each agent in all previous periods.

LLMs are instructed through prompts. The *user prompt* sets the task or query that the user wants the model to respond to, and it can change with each interaction. In addition to the user prompt, LLMs can also be instructed through their *system prompt*, which sets the context, behavior, knowledge base, and role for the model. We use the system prompt to provide the general instructions of the experiment (corresponding to the written instructions handed out to human participants) and the user prompt to provide updates throughout the trading periods and request trading actions. The full text of all prompts is provided in Appendix A.

Timeline for each round in the AI laboratory:

1. A computer selected the asset’s fundamental value, as in the human experiment.
2. We make an API call to an LLM, using the instructions of the experiment as the system prompt, see Prompt 1. The user prompt requests the model to provide a trading action (buy, sell, no trade) given each signal (blue and white) and the current asset price, along with its reasoning for each action. For trading rounds $t > 1$, the user prompt also provides, for each agent, the history of executed trades, a notification if that agent was chosen to act in the previous round, and the history of actions and reasoning of that agent. The user prompt is provided in Prompt 2.
3. The trader selected trade and the realized signal are chosen by a computer, as in the human experiment.
4. The price for the next round is computed given the selected trader’s action for the realized signal.
5. Steps 2-4 are repeated for eight trading periods total.

This timeline is summarized in flow diagrams in Figure 1 for each treatment.

Each experiment (i.e., four sessions of eight trading rounds) is repeated across different LLMs. To maintain comparability across experiments, we seed the randomness such that the realized asset value, realized signals, and the sequence of selected traders are identical across the experiments. Following Cipriani and Guarino (2009), we assume that an information event always happens, even in Treatment II, where the theoretical probability of an information event is less than one.

Optimal AI laboratory To establish a baseline, our main AI laboratory does not instruct the LLMs about what constitutes optimal decision making. In contrast, real-world adoptions of AI in financial

market decision making likely attempts to guide the models in optimal behavior to the largest extent possible. We therefore also explore an optimal version of the AI laboratory, where we explicitly tell the models when herding is optimal according to theory through the user prompt. This prompt is given in Prompt 3.

5. Results

This section presents the results from conducting the Cipriani and Guarino (2009) experiments in the AI laboratories. We mainly focus on the baseline AI laboratory in which we do not include any guidance on optimal decision making in the prompt. AI agent decisions are presented alongside human decisions from the original experiment. First, we consider the case without event uncertainty (Treatment I), then we introduce event uncertainty (Treatment II), and finally we prevent price updating (Treatment III). We also analyze the descriptions of reasoning provided by the LLMs to understand the decisions of AI agents. Unfortunately, we do not have human counterparts for these insights.

These baseline results establish patterns of decision making in general-purpose technology without further optimization. In real-world applications, professional traders would likely consult tools that have been fine-tuned to maximize profits according to some risk-profile. To examine herding behavior in such optimized AI agents, we consider an AI laboratory where the LLMs are informed about optimal decisions through the prompt. Finally, we consider robustness to the model temperature parameter.

5.1. Without event uncertainty

We begin by discussing results obtained with the parameterization in Treatment I, where there is no model uncertainty. The theoretical model predicts that traders should always trade according to their private signals, which precludes the formation of cascades.

Table 1 shows the frequency of the different behaviors averaged across all sessions and trading periods. The “Human” column recites the results from the human laboratory reported in Cipriani and Guarino (2009). The “AI” column represents the average results across all considered LLMs. With Treatment I, reported in panel (a), AI agents exhibit more rational behavior, i.e., buy on a “good” signal and sell on a “bad signal,” (61%) compared to humans (46%). This result is largely driven by the Claude 3.7 and Llama 3 models, which generate rational responses for a vast majority of sessions and trading periods. In contrast, the Claude 3.5 and Nova Pro models have fewer rational responses, but a majority of responses that are partially rational, i.e., follow the rational response on one signal but decide to not trade on the other signal. As a result, the share of rational and partial rational responses in the AI laboratory far exceed that observed in the human laboratory (90% for AI versus 65% for humans). It is worth noticing

that while humans make mistakes (in 3.40% of the total decisions), no erroneous decisions were made in the AI laboratory.

Information cascades occur in less than 10% of the decisions in the AI laboratory, less than one third of the frequency of information cascades in the human laboratory. Cascade trading behavior is mostly driven by the Claude models, while Nova Pro is the only model that generates no-trade cascades.

We can gauge the nature of these cascades when there is a trade imbalance, i.e., a difference between the number of sell and buy orders in the trading history. Information cascades represent herding if the cascade follows the market, i.e., the majority action in the trading history, and contrarian behavior if the cascade goes against the dominant action in history. Table 1 also shows the decomposition of cascade trading into optimal and suboptimal herding, contrarian, and undetermined behavior. Furthermore, the table reports the fraction of decisions where herding is optimal (which is equal to zero percent in Treatment I). The results show that the trading cascades are fully attributed to contrarian behavior. While the human experiment does identify some herding, it is also the case here that contrarian behavior is dominating.¹³ Neither herding nor contrarian behaviors, however, are predicted to be optimal by the theoretical model.

Why do LLMs engage in contrarian behavior although theory predicts that this type of decision is never optimal? One explanation is that AI agents fail to incorporate trading history into their expectations of the asset's value. Without this information, the agent will value the asset at the price of 70 ($0.7 * 100 + 0.3 * 0$) given a white signal 30 ($0.3 * 100 + 0.7 * 0$) given a blue signal. In contrast, the market maker *does* take trading history into account when updating the price. For example, after a sufficient number of buy orders, the price will increase above 70. In such case, an agent who ignores the trading history will sell regardless of the signal, as the expected value is lower than the price given both signals, hence engage in contrarian behavior. Analyzing the reasoning provided by the LLMs confirms this explanation, see Section 5.4.

5.2. With event uncertainty

With Treatment II, where there is uncertainty about whether an information event has occurred, herding can be optimal in the theoretical model. Contrarian behavior, however, is never optimal.

Panel (b) of Table 1 shows that none of the AI agents decide to herd during the experiment as cascade trading behavior is non-existent. The AI agents also do not engage in contrarian behavior, consistent

¹³ Cipriani and Guarino (2009) reports the decomposition of cascade trading behavior by trade imbalance. But, as the distribution of trade imbalance across trading periods is unknown, we cannot infer the exact decomposition in Table 1. We overcome this limitation by simulating the trading imbalance using 10,000 independent iterations, which allows us to simulate the decomposition of cascade trading behavior.

with theory. In fact, most of the decisions are driven by private information (97%), which far exceed this share of decisions among human participants (51%). The table also reports the percentage of times where herding is optimal. In both laboratories (human and AI), herding is optimal in a little more than one third of the times. The LLMs overlook these opportunities in their focus on the information contained by the private signals. We note that humans also do not exploit all of these opportunities as they engage in optimal herding in just 4% of their decisions, which accounts for just one third of their cascade trading behavior.

5.3. The impact of price updating

The difference between Treatment I and II lies primarily in the price updating rule. Specifically, not knowing whether an information event has occurred, the market maker updates the price more conservatively in Treatment II. Figure 2 illustrates the price dynamics for each trading round, averaged across the experiments for each LLM. Each line is a session. The figure shows that under Treatment I, the price moves away from the initial price of 50. In two of the sessions, the price is close to zero or 75 after eight periods. In contrast, under Treatment II, the price stays close to the initial price of 50 throughout all trading periods.

At the extreme, when the price does not update at all (Treatment III), see panel (c) of Table 1, all models make decisions relying only on private information in practically all sessions and trading rounds (more than 99%). This result is obtained despite the fact that in around one third of the decisions herding would have been optimal (when the trade imbalance exceeds two).

5.4. Analyzing LLM reasoning for trading decisions

While we do not know the reasoning behind the decisions made in the human laboratory, we asked the LLMs as part of the user prompt to give reasoning for their decisions. Analyzing these reasoning paragraphs may shed further light on the decision making process in each of the models. It is not guaranteed that the reasoning provided by LLMs reflects the actual reasoning applied by these models to arrive at their decisions. However, as we show below, there is consistency between reasoning and decisions.

We examine the lines of reasoning using both LLMs and the LDA topic model. For the LLM analysis, we use the Claude 3.7 model—the most advanced among our models—to characterize each passage of reasoning. Specifically, for each passage, we ask the LLM to read the reasoning and answer the following five questions:

Question 1: Is the trader comparing the price to the expected fundamental value of the asset?
(True/False)

Question 2: Is the expected value computed using only the signal accuracy and the signal, e.g., $0.7 * 100 + 0 * 0.3 = 70$ or $0.7 * 0 + 0.3 * 100 = 30$? (True/False)

Question 3: Does the trader consider the market trend or the trading history in their reasoning? (True/False)

Question 4: How does the trader characterize the attractiveness of the investment (very attractive, attractive, reasonable, less attractive, no incentive)?

Question 5: On a scale from 0-100 (where 100 represents purely emotional and 0 represents purely rational or logical), how much is the investor driven by emotions in their assessment?

Table 2(a) shows the average responses to each of these questions across all LLMs for all reasoning passages belonging to each treatment. A positive response to the first question is a necessary conditional for rational decision making. Indeed, the results shows that this condition is satisfied for basically all passages. Question 2 seeks to understand if the AI agents evaluate expected values given only the signal, ignoring trading history, as conjectured from the distribution of decisions. For all treatments, this happens in nearly two thirds of the decisions confirming our explanations for our results stated earlier. For example, in the seventh trading period of the second session, two participants chose to buy on both signals at the price of 15.52, forming a cascade. One of the agents gave the following reasoning for buying on a blue (“bad”) signal:

“Even with a Blue signal, the expected value is 30 (30% chance of 100, 70% chance of 0). The current price of 15.52 still below this expected value, buying remains profitable.”

This argument disregards that the trading history (in this case {buy, no trade, sell, sell, sell, sell}) and the total price decrease from the initial price of 50 to 15 indicate that the asset value is zero, assuming that other participants followed rational responses such that the trading history reflects private information from previous periods.

Question 3 tackles the same question from a different angle. We find that 10-24% of the decisions are (at least partly) based on the market trend or the trading history. Hence, consistent with answers to the second question, a majority of the AI agents do not consider the trading history when forming expectations.

Disregarding the accumulation of private information in the pricing of the asset also explains the large share of partially rational decisions, i.e., why the model would decide to not trade on one of the signals. For example, in the fourth round of the fourth session, all of the Claude 3.5 agents decided to not trade at the price of 30 given a blue (“bad”) signal because, as one of the agents put it:

“With a Blue signal, the expected value is 30 (30% chance of 100, 70% chance of 0). Since this matches the current price, there’s no clear advantage to trading.”

If the model had taken into account the fact that the majority action in the trading history was to sell the asset, it may have assigned a higher probability of a low asset value than the signal accuracy of 70%, arriving at a lower expected value and consequently decided to buy the asset.

In the fourth question, we examine the certainty with which AI agents make their decisions. More than half of the decisions are evaluated as either very attractive or attractive and less than one third is deemed unattractive. These results tell us that the results do not represent LLMs making enforced decisions with potential arbitrary decision outcomes: AI agents generally deem that there is an opportunity to make reasonable profits by engaging in trading. The absence of no-trading cascade decisions in Table 1 supports this conclusion as well.

Finally, we ask Claude 3.7 to evaluate the degree of emotional content in the reasoning on a scale from 0-100. Consistent with the answers to the first three questions, these scores are generally low with averages about 5%, top deciles of 15-20%, and medians of zero.

These results are similar for reasoning passages generated by all four LLMs, with the exception of the Llama 3 model. For this reason, analysis of Llama 3 reasoning is provided separately in Table 2(b). This model appears to use more “judgment” or “emotion” in its reasoning. For example, the scores from Question 5 measuring the degree of emotion average at 13-17% for the Llama 3 reasoning passages. Also, responses to Questions 2-3 indicate that Llama 3 does not reason using the expected value conditional on the signal alone, but includes trading history and market trends to a greater extent than the other models. For example, a representative reasoning for Llama agents (from the first round of session one):

“Since the White signal is more likely to occur when the asset value is 100, I believe the asset is more likely to be worth 100 than 0, so I’m willing to buy at a price of 50. This is a good deal for me if the asset is indeed worth 100.”

Another Llama agent from the same session considers the potential of the market maker inflating the price of the asset (from the third round of session one):

“Conversely, the Blue signal suggests that the fundamental value might be 0. Selling at 84.48 seems like a good opportunity to get rid of a potentially worthless asset, especially since the market maker’s updated price might be overestimating the asset’s value.”

We confirm these results using a more traditional approach to topic analysis, namely the LDA topic model, in Electronic Appendix.

5.5. Optimal AI agents

By not taking into account the cumulation of private information in the trading history, the AI agents avoid suboptimal herding behavior in Treatment I, but also overlook potential optimal herding in Treatment II and III. In contrast, trading cascades arise in the human laboratory both when such are strictly suboptimal as in Treatment I and when they can be optimal in Treatment II.

Real-world integrations of AI in investing will likely involve fine-tuning the models to behave as optimal and profit-maximizing as possible. This is obviously a difficult task as it is not known a priori which decisions are optimal. In contrast, in our controlled experimental setting, we know from theory which decision is optimal in expectation, and we can prompt the LLMs directly with this information. We thus implement the experiment in an *optimal AI laboratory*, which is similar to the AI laboratory described in Section 4.2 except that we explicitly prompt the LLMs when herding is optimal. The user prompt for this exercise is provided in Prompt 3. Since the human experiment did not provide human participants with these detailed guides on optimal behavior, we cannot compare the decisions of optimal AI agents with human decisions.

Table 3 shows the results. In Treatment I, the optimal AI agents engage less frequently in cascade trading behavior, which is reduced to 3.5% of the decisions as compared with 9.4% in the baseline results in Table 1. In Treatment II, herding is optimal in 81.51% of the decisions on average across LLMs, and the optimal AI agents herd in 47.43% of the decisions. There is no suboptimal herding behavior, but the optimal AI agents do make contrarian decisions and involve in cascade trading when the trade imbalance is zero. Finally, results for Treatment III show that optimal AI agents exploit most of the optimal herding opportunities (on average 44.36% out of 50.90%). There is, however, some suboptimal herding.

In summary, by explicitly including the optimal behavior in the prompt given to the LLMs, we urge the models to herd when optimal. But, we note that the models do not exploit all of the opportunities for optimal herding and that the optimal AI agents make more suboptimal cascade trading decisions compared with the baseline AI agents. Attempting to fine-tune LLMs to behave optimally can thus have unintended consequences and, in turn, financial stability implications. Despite these cases, optimal AI agents generally have higher expected payoffs than the baseline AI agents, as shown in Table 4. The difference in expected payoff is particularly pronounced in Treatment II, where the baseline AI agents earn an average of 3.8 lire and the optimal AI agents earn 15 lire on average. In other words, baseline AI agents are punished for avoiding to herd when optimal.¹⁴

¹⁴ In Treatment II and III, where the price stays near or exactly at 50, the consequence of making a bad decision (buying an asset worth zero or selling an asset worth 100) is larger than in Treatment I, where the price moves towards the fundamental value.

5.6. Robustness to model temperature

Temperature is a hyperparameter to LLMs controlling how the models predict the next token in a sequence. With a lower temperature, the model is more likely to choose the most probable next token, resulting in more deterministic and less creative responses. Higher temperatures flatten the probability distribution of the next token leading to more variation in the responses. For the baseline results, we applied a medium temperature of 0.7 for all models except Claude 3.7, which we apply in “extended thinking mode” which only supports a temperature of 1.0. Electronic Appendix shows that the temperature setting only has minor impact on the decisions of AI agents in the experiments.

5.7. Robustness to alternative signal codings

Theoretically, it does not matter if the “good” signal is white and the “bad” signal is blue in the experimental design. However, it may matter in practice. For example, Bazley et al. (2021) show that the perception of color for visualizing financial data influences individuals’ risk preferences, expectations of future stock returns, and trading decisions. Specifically, the color red has been associated with higher probabilities assigned to loss outcomes (Kliger & Gilad, 2012) and more risk averse behavior (Gnambs et al., 2015). Testing different signal colors in the AI laboratory therefore serves as a test of whether LLMs work purely as algorithmic robots (for whom the labeling of signals is irrelevant) or are contaminated by human bias (whose actions are impacted by the choice of signal labels).

Simply asking LLMs to associate financial market conditions with a color-coded signal reveals that the models perceive white and blue as neutral signals indicating stable market conditions.¹⁵ In contrast, the models associate green and red with market movements, bullish and bearish, respectively.¹⁶

We test two alternative versions of signals. The first variation tests an experiment where a “good” signal is represented by the color green and a “bad” signal is represented by the color red. This variation is arguably more charged with meaning or connotation than the baseline of white/blue signals, but the alignment of green with “good” and red with “bad” adheres to typical Western color associations. In the second variation, we reverse the labeling such that a “good” signal is represented by red and a “bad” signal by blue.

Results are shown in Table 5. Due to significant variations in outcomes across different models, we present the results for each individual LLM separately in Electronic Appendix. Using green/red in place of white/blue generally does not impact the model-averaged results, in any of the treatments. In contrast,

15 Llama 3 consider blue as a bullish or positive market signal, which interestingly does not impair with its decisions in the baseline experiments, where blue is used as a “bad” signal.

16 These responses are documented in Electronic Appendix.

when we invert the conventional color associations by using red to indicate “good” and green to signify “bad”, we observe a dramatic shift in the results. First, on average, the models generate errors, i.e., decisions to sell given a “good” signal and buy given a “bad” signal, in one quarter of all decisions in both treatments. This result is driven by Claude 3.5 for which all decisions are erroneous under this color scheme.

Second, we observe more cascade trading decisions under Treatment II and III. These decisions are driven by the Llama 3 model, and represent herding or cascade trading under zero trade imbalance (cannot be determined as herding or contrarian behavior). Inspecting the model’s explanations reveals that the model understands that red is a “good” signal, but at the same time associates green with “good,” or remains optimistic even with a “bad” signal. For example, in the first round of session four, an agent gives the following reasoning for choosing to buy given a green (“bad”) signal:

“Although the Green signal is more likely to occur when the fundamental value is 0, there is still a 30% chance it’s 100. Buying at 50 is still a good deal, considering the expected value is 50.”

The model also associates the signals with the correct interpretation in the cases where it decides to follow the herd. For example, in the seventh round of the fourth session, an agent chooses to buy at the price of 99 on both signals, and provides the following explanation for the buy-decision on the green (“bad”) signal:

“Although the probability of the asset value being 100 given a Green signal is lower than with a Red signal, it is still quite high. The increasing market price and previous buy decisions suggest that the market is moving towards a higher valuation, making it a good buy even with a Green signal.”

The Llama model thus seems to carry over its interpretation of green as a positive signal, despite the clear instructions that green signals a high probability of the asset being worthless. The Nova Pro model mainly rely on private information under both color schemes, and thus behaves as one would expect from algorithm-driven intelligence.

These results could have important implications for financial stability. While AI agents appear mainly rely on private information when information arrives in ways that conform with expectations (e.g., green implies a positive signal), they can generate irrational or even erroneous outcomes if the meaning of a type of signal changes. This could be an important factor if a shock produces responses that are unexpected given past experience. More broadly, these findings suggest that AI agents are not purely algorithmic decision makers who objectively process given information, but are susceptible to preconceived biases.

These results also demonstrate that the choice of labels in experimental design can substantially influence outcomes, particularly when these labels contradict intuitive associations. We conjecture that

similar effects would likely be observed in the human laboratory.

Interestingly, the most recent and advanced model in terms of reasoning capability—the Claude 3.7 model—generates similar results regardless of the signal. Thus, as LLMs improve, this risk may be partly mitigated.

6. Implications for financial stability

Our results suggest that AI agents exhibit less herd behavior at the extensive margin compared to human financial professionals. The reduced tendency of AI to herd has significant implications for financial stability as generative AI gains traction in market decision making. This section first presents simulations based on the experiments to illustrate how behavior of AI agents and humans impact the convergence of prices to fundamental values. Then, we present four conjectures of how our results could impact financial stability.

6.1. Simulating price convergence

To illustrate potential financial stability implications, we use the experimental results to simulate pricing errors in a more realistic setting where human and AI-powered traders interact over many periods.

First, we simulate the market in our experimental setting over 100 trading periods in a theoretical world where all traders exhibit either rational, herding, contrarian, or optimal behavior.¹⁷ We adopt the parametrization in Treatment II, where herding is sometimes optimal. Figure 3 shows that when decision-making is driven by private information (rational behavior), errors quickly converge to zero. This result emerges despite the fact that this rational behavior is not always optimal, and is thus similar to results in Gode and Sunder (1993) showing that such convergence can occur even when traders do not attempt to maximize profits or pay attention to market events. In contrast, herding and contrarian behavior prevents the price from converging to the fundamental value, on average. When traders always act optimally to maximize their expected profits, pricing errors initially adjust towards zero more rapidly than errors under fully rational behavior. But on average, errors do not converge to zero because actions that are optimal in expectation may turn out to be suboptimal, e.g., if an information cascade is initiated based on a wrong signal.¹⁸ Based on the median pricing error, where such cases are not represented, information cascades ensure faster convergence than rational behavior.

These results illustrate that both optimal and suboptimal herding can lead to financial fragility, though

¹⁷ Simulations are conducted with 1000 independent iterations.

¹⁸ Indeed, pricing errors converges to lower numbers as the signal accuracy is increased.

through different mechanisms. Figure 4 elaborates on such channels. Optimal herding can accelerate information aggregation and enhance efficiency by incorporating dispersed knowledge into prices (Chamley, 2003). For example, withdrawals from a genuinely insolvent bank reflect individually rational and collectively justified behavior. Such market discipline may improve financial stability if the market is able to distinguish between insolvent banks and solvent banks. Nevertheless, even optimal herding may generate short-term volatility by accelerating price corrections, which can reverse quickly if new information arrives. Short-term volatility can also arise from concentrating trading flows, potentially exacerbating fire sales and liquidity strains. Finally, optimal herding can overshoot or undershoot fundamental values, or transmit to other markets, triggering suboptimal herding with adverse consequences (Cipriani & Guarino, 2008). These dynamics create a paradox: actions that optimize individual utility can simultaneously undermine market efficiency and stability. Financial stability may therefore be enhanced when investors act on their private signals rather than engage in behavior that is individually optimal but systemically destabilizing.

Next, we simulate 100 trading periods, where traders all follow the distribution of decisions of respectively AI agents, humans, and Optimal AI agents uncovered in our experiments. We also simulate markets where 50% of the traders are AI agents and 50% of the traders are human, which captures human-AI interactions given the fixed distribution of decisions. For feasibility, we assume that the distribution of trading decisions uncovered from the experiments remain constant.¹⁹ Details are provided in Appendix B.

Figure 5 shows the average and median absolute pricing error for each trading period in Treatment I (panel a) and Treatment II (panel b). In Treatment I, where herding is never optimal, AI agents ensure that prices converge quickly towards the fundamental value of the asset, even when half of the traders are humans. When humans trade alone, pricing errors persist through more trading periods, but do eventually converge to zero. Similar dynamics are observed in Treatment II. In this treatment, however, average pricing errors do not converge to zero among Optimal AI traders, because of their tendency to engage in optimal herding. Similar to the pricing errors under optimal behavior in Figure 3, this result is driven by a few outlier cases where an information cascade was initiated based on a wrong signal. Indeed, the median pricing error does converge to zero, and at a faster pace than the AI and human agents.²⁰

19 Electronic Appendix shows that prolonging the experiment from eight to 20 trading rounds does not substantially change the distribution of decisions among AI agents.

20 Appendix B shows that the upper quartile of the 100th period pricing error of Optimal AI agents is just 0.13 while the maximum is 99.99.

6.2. Four conjectures

Our controlled experimental approach isolates behavioral mechanisms of AI-powered trading that may carry implications for financial stability. First, our findings that AI avoids extensive margin herding suggest that widespread AI adoption could reduce the frequency of coordination failures that lead to market freezes and extreme market movements. As AI systems increasingly influence trading decisions, either directly through algorithmic trading or indirectly by advising human investors, markets may become less prone to self-reinforcing cycles that drive prices away from fundamentals.

Second, if AI can be successfully instructed to herd when herding is optimal, financial stability implications are more nuanced. Optimal herding may lead to faster price discovery and correction, which could uncover existing vulnerabilities more quickly, potentially enabling earlier regulatory or market responses before problems become systemic. However, this same property could increase short-term volatility and lead to more abrupt market adjustments. Also, an increased tendency to herd (even when optimal) can also drive prices away from fundamentals, because signals and interpretations of information could be wrong; for example, in our experiment, signals are wrong with 30% probability. Importantly, we emphasize that the trading decisions of Optimal AI agents are not comparable with the human decisions as the human experiment was not repeated in this experimental setup. Also, inducing AI agents to engage in optimal herding required very detailed guidance, which is likely challenging to achieve beyond our controlled laboratory setting.

Third, AI's stronger tendency toward relying on private information in their trading decisions may diversify market participant reactions to new information. Rather than all participants interpreting and acting on information in similar ways, AI might introduce greater heterogeneity in responses, potentially reducing correlation in market movements. On the other hand, AI may also be adopted to help interpreting information, corresponding in our experiment to labeling signals as either "good" or "bad" (white or blue). In this case, AI may cause more homogeneity in the signals that traders receive, which would increase market correlation.

Fourth, the results from re-labeling signals in the experiment reveal that AI is not perfectly rational. When signals in experiments were deliberately labeled counter-intuitively, LLMs produced few responses driven by private information, suggesting they have inherited elements of human intuition and bias. This hybrid nature of AI decision-making—less driven by market trends than humans but not purely rational—creates additional complexity in predicting how widespread AI adoption might impact financial stability. While AI may reduce certain human biases, it introduces its own form of imperfect rationality that must be accounted for in stability assessments.

It is important to note that these implications are speculative and based on experimental results. The actual impact of AI on financial stability will depend on numerous factors, including the extent of AI

adoption, the specific models used, regulatory responses, and how AI systems evolve over time, which may incorporate more sophisticated agentic AI frameworks. In addition, the interaction between human and AI traders becomes crucial, as their combined behavior could either amplify or dampen market movements in unpredictable ways. This shift may necessitate new tools and approaches for regulatory oversight, including AI-specific stress tests or new forms of market surveillance. Furthermore, the long-term implications of AI decision-making on market stability, including potential unforeseen consequences, remain an important area for further research. Notably, traditional measures of market sentiment, which often rely on human emotions and behaviors, may need to be reconsidered. With increased AI involvement, new methods may be needed to gauge market sentiment and predict potential instabilities, as the emotional drivers of market behavior could shift significantly.

7. Conclusion

This study establishes stylized facts about AI trading behavior under uncertainty, offering micro-founded insights into the potential impact of AI on financial stability. We compare the decision making behavior of AI agents with that of human financial professionals in a controlled experimental setting studying herd behavior at the extensive margin. Our findings show that AI agents rely more on private information and demonstrate less propensity for information cascades compared to their human counterparts. In fact, AI agents avoid herding even under conditions when herding is optimal, and they do not exploit all optimal herding strategies even when explicitly advised when herding is optimal.

Our work also reveals a potential paradox that may occur in a world of AI-powered financial decision making. AI trading decisions are different from human decision-making, yet in certain situations remain tethered to human biases. When signal labels contradict human conditioning (red signaling “good” rather than “bad”), rational behavior collapses as AI appears to have absorbed the irrational residue of their training data. In other words, AI agents appear to be averse to animal spirits but, at the same time, also lack complete algorithmic discipline.

The financial stability implications cut both ways. Reduced herding behavior could dampen extreme market movements driven by irrational exuberance, which can lead to greater overall market stability, as our simulations demonstrate faster price convergence toward fundamentals. However, optimal AI agents—explicitly guided to maximize profits—do engage in cascade trading, suggesting that more fine-tuned deployment could lead to herding dynamics, potentially increasing market volatility, but through superior or optimal coordination. The hybrid nature of AI decision-making also introduces a new source of fragility, for example through AI agents that behave algorithmically under expected or neutral conditions but fail when information arrives in unexpected or counterintuitive forms.

As financial markets undergo their most profound transformation since electronic trading, this paper provides an intellectual framework for monitoring how AI integration in the financial system enhances or threatens financial stability. More AI integration in financial transactions may introduce new forms of financial stability considerations that existing regulatory frameworks may not be able to anticipate. Indeed, understanding AI's behavioral foundations becomes essential for monitoring financial stability in an age where AI will likely play a significantly larger role in the financial system.

Statement on the Use of AI Tools

In preparing this manuscript, the authors used Claude 3.7 Sonnet (Anthropic), Claude 3.5 Sonnet (Anthropic), Llama 3 (Meta), and Amazon Nova Pro (Amazon) were used as part of the analysis within the AI laboratory set-up described in Section 4. In preparing the manuscript, the authors additionally used Claude Sonnet 4.5 (Anthropic) to improve the clarity and readability of draft text and to provide feedback on manuscript drafts, which the authors evaluated alongside comments from human reviewers and discussants. All AI-assisted content, including text and figures, was critically reviewed, edited, and verified by the authors, who take full responsibility for the accuracy, originality, and integrity of the published work.

References

- Anand, K., Kazinnik, S., Leonello, A., & Panetti, E. (2025). Ex Machina: Financial Stability in the Age of Artificial Intelligence. *SSRN Working Paper*.
- Angeletos, G.-M., Collard, F., & Dellas, H. (2018). Quantifying Confidence. *Econometrica*, 86(5), 1689–1726.
- Argyle, L. P., Busby, E. C., Fulda, N., Gubler, J. R., Rytting, C., & Wingate, D. (2023). Out of One, Many: Using Language Models to Simulate Human Samples. *Political Analysis*, 31(3).
- Avery, C., & Zemsky, P. (1998). Multidimensional Uncertainty and Herd Behavior in Financial Markets. *American Economic Review*, 88(4), 724–748.
- Banerjee, A. V. (1992). A Simple Model of Herd Behavior. *The Quarterly Journal of Economics*, 107(3), 797–817.
- Bank of England Financial Policy Committee. (2025). *Financial Stability in Focus: Artificial Intelligence in the Financial System* (tech. rep.). Bank of England.
- Bazley, W. J., Cronqvist, H., & Mormann, M. (2021). Visual Finance: The Pervasive Effects of Red on Investor Behavior. *Management Science*, 67(9), 5616–5641.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades. *Journal of Political Economy*, 100(5), 992–1026.
- Bikhchandani, S., & Sharma, S. (2001). Herd Behavior in Financial Markets. *IMF Staff Papers*, 47(3), 279–310.
- Bini, P., Cong, L. W., Huang, X., & Jin, L. J. (2025). Behavioral Economics of AI: LLM Biases and Corrections. *SSRN Working Paper*.
- Bybee, J. L. (2023). The Ghost in the Machine: Generating Beliefs with Large Language Models. *SSRN Working Paper*.
- Cartea, Á., Chang, P., Chen, N., & Zhong, M. (2026). AI Bubbles with Large Language Models. *SSRN Working Paper*.
- Chamley, C. P. (2003). *Rational Herds*. Cambridge University Press.
- Chen, Y., Liu, T. X., Shan, Y., & Zhong, S. (2023). The Emergence of Economic Rationality of GPT. *Proceedings of the National Academy of Sciences (PNAS)*, 120(51), 1–9.

- Cheng, Q., Lin, P., & Zhao, Y. (2025). Does generative AI facilitate investor Trading? Early evidence from ChatGPT outages. *Journal of Accounting and Economics*, 80(2), 101821.
- Cipriani, M., & Guarino, A. (2005). Herd Behavior in a Laboratory Financial Market. *American Economic Review*, 95(5), 1427–1443.
- Cipriani, M., & Guarino, A. (2008). Herd Behavior and Contagion in Financial Markets. *The B.E. Journal of Theoretical Economics*, 8(1), 1–54.
- Cipriani, M., & Guarino, A. (2009). Herd Behavior in Financial Markets: An Experiment With Financial Market Professionals. *Journal of the European Economic Association*, 7(1), 206–233.
- Danielsson, J., & Uthemann, A. (2024). AI Financial Crises. *VoxEU.org*.
- del Rio-Chanona, R. M., Pangallo, M., & Hommes, C. (2025). Can Generative AI Agents Behave Like Humans? Evidence From Laboratory Market Experiments. *arXiv Preprint arXiv:2505.07457v1*.
- Dou, W. W., Goldstein, I., & Ji, Y. (2025). AI-Powered Trading, Algorithmic Collusion, and Price Efficiency. *SSRN Working Paper*.
- Drehmann, M., Oechssler, J., & Roeder, A. (2005). Herding and Contrarian Behavior in Financial Markets: An Internet Experiment. *American Economic Review*, 95(5), 1403–1426.
- Fedyk, A., Kakhbod, A., Li, P., & Malmendier, U. (2024). AI and Perception Biases in Investments: An Experimental Study. *SSRN Working Paper*.
- Financial Stability Board. (2024). *The Financial Stability Implications of Artificial Intelligence* (tech. rep.). Financial Stability Board.
- Galariotis, E. C., Krokida, S.-I., & Spyrou, S. I. (2016). Herd Behavior and Equity Market Liquidity: Evidence From Major Markets. *International Review of Financial Analysis*, 48, 140–149.
- Gnambs, T., Appel, M., & Oeberst, A. (2015). Red Color and Risk-Taking Behavior in Online Environments. *PLOS One*.
- Gode, D. K., & Sunder, S. (1993). Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, 101(1), 119–137.
- Grauwe, P. D. (2012). *Lectures on Behavioral Macroeconomics*. Princeton University Press.

- Hansen, A. L., Horton, J. J., Kazinnik, S., Puzzello, D., & Zarifhonarvar, A. (2025). Simulating the Survey of Professional Forecasters. *SSRN Working Paper*.
- Hartley, J., Jolevski, F., Melo, V., & Moore, B. (2024). The Labor Market Effects of Generative Artificial Intelligence. *SSRN Working Paper*.
- Hayes, W. M., Yax, N., & Palminteri, S. (2024). Relative Value Biases in Large Language Models. *arXiv Preprint arXiv:2401.14530*.
- Henning, T., Ojha, S. M., Spoon, R., Han, J., & Camerer, C. F. (2025). LLM Trading: Analysis of LLM Agent Behavior in Experimental Asset Markets. *Proceedings of the Thirteenth International Conference on Learning Representations (ICLR)*.
- Horton, J. J. (2023). Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus? *NBER Working Paper, No. 31122*.
- Hsieh, S.-F., Chan, C.-Y., & Wang, M.-C. (2020). Retail Investor Attention and Herding Behavior. *Journal of Empirical Finance*, 59, 109–132.
- Hua, W., Liu, O., Li, L., Amayuelas, A., Chen, J., Jiang, L., Jin, M., Fan, L., Sun, F., Wang, W., Wang, X., & Zhang, Y. (2024). Game-theoretic LLM: Agent Workflow for Negotiation Games. *arXiv Preprint arXiv:2411.05990*.
- International Monetary Fund. (2024). *Global Financial Stability Report* (tech. rep.).
- Jha, M., Qian, J., Weber, M., & Yang, B. (2024). *ChatGPT and Corporate Policies* (tech. rep.). National Bureau of Economic Research.
- Jia, J., Yuan, Z., Pan, J., McNamara, P. E., & Chen, D. (2024). Decision-Making Behavior Evaluation Framework for LLMs under Uncertain Context. *arXiv Preprint arXiv:2406.05972*.
- Jiang, H., Zhang, X., Cao, X., Kabbara, J., & Roy, D. (2023). PersonaLLM: Investigating the Ability of GPT-3.5 to Express Personality Traits and Gender Differences. *arXiv Preprint arXiv:2305.02547*.
- Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., & Mullainathan, S. (2018). Human Decisions and Machine Predictions. *The Quarterly Journal of Economics*, 133(1), 237–293.
- Kliger, D., & Gilad, D. (2012). Red Light, Green Light: Color Priming in Financial Decisions. *The Journal of Socio-Economics*, 41(5), 738–745.

- Koralus, P., & Wang-Maścianica, V. (2023). Humans In Humans Out: On GPT Converging Toward Common Sense in Both Success and Failure. *arXiv Preprint arXiv:2303.17276*.
- Lee, H., Seo, J., Park, S., Lee, J., Ahn, W., Choi, C., Lopez-Lira, A., & Lee, Y. (2025). Your AI, Not Your View: The Bias of LLMs in Investment Analysis. *arXiv Preprint arXiv:2507.20957*.
- Li, C., Wang, H., Jiang, S., & Gu, B. (2024). The Effect of AI-Enabled Credit Scoring on Financial Inclusion: Evidence from an Underserved Population of over One Million. *MIS Quarterly*, 48(4), 1803–1834.
- Liu, R., Geng, J., Peterson, J. C., Sucholutsky, I., & Griffiths, T. L. (2025). Large Language Models Assume People are More Rational than We Really Are. *Proceedings of the Thirteenth International Conference on Learning Representations (ICLR)*.
- Lopez-Lira, A. (2025). Can Large Language Models Trade? Testing Financial Theories with LLM Agents in Market Simulations. *arXiv Preprint arXiv:2504.10789*.
- McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L. S., & Zolas, N. (2023, October). "AI Adoption in America: Who, What, and Where" (Working Paper No. 31788). National Bureau of Economic Research.
- Muntoni, M. (2025). AI Expectations and Market Stability. *SSRN Working Paper*.
- Ross, J., Kim, Y., & Lo, A. W. (2024). LLM economicus? Mapping the Behavioral Biases of LLMs via Utility Theory. *arXiv Preprint arXiv:2408.02784*.
- Scharfstein, D. S., & Stein, J. C. (1990). Herd Behavior and Investment. *The American Economic Review*, 80(3), 465–479.
- Shiller, R. J. (2005). *Irrational Exuberance* (2nd ed.). Princeton University Press.
- Wang, S., Zhang, S., Zhang, J., Hu, R., Li, X., Zhang, T., Li, J., Wu, F., Wang, G., & Hovy, E. (2024). Reinforcement Learning Enhanced LLMs: A Survey. *arXiv Preprint arXiv:2412.10400v3*.
- Zarifhonorvar, A. (2024). Experimental Evidence on Large Language Models. *Available at SSRN*, 4825076.
- Zhu, J.-Q., & Griffiths, T. L. (2024). Incoherent Probability Judgments in Large Language Models. *arXiv Preprint arXiv:2401.16646*.

Figures and Tables

Figure 1: Flow diagrams of experiments

The figure shows diagrams of the order of events for each session of the experiments under (a) Treatment I (without event uncertainty) and III (without price updating), and (b) Treatment II (with event uncertainty). The experiment is an adoption of Cipriani and Guarino (2005, 2009) and is based on the Avery and Zemsky (1998) model.

(a) Treatment I and III




Figure 1: Flow diagrams of experiments (*continued*)

(b) Treatment II




Figure 2: Price dynamics

The figure shows the price dynamics across trading periods for each treatment, averaged across LLMs in (a) Treatment I (without event uncertainty) and (b) Treatment II (with event uncertainty). Each line represent one of the four independent sessions. Following a buy or sell order, the price is updated by a Bayesian market maker given the trading history.




Figure 3: Simulated pricing errors: Theoretical behaviors

The figure shows the average (left) and median (right) absolute pricing errors across trading periods simulated in the experimental model under Treatment II with 1000 independent iterations. The lines show four different behavioral assumptions: all traders are rational, i.e., trade based on their private signals; all traders herd, i.e., follow the market trend; all traders are contrarian, i.e., trade against the market trend, and all traders act optimally to maximize expected profits (either rational or optimal herding behavior).




Figure 4: Herding behavior and financial stability

The diagram shows how herding can lead to a financial stability event both when optimal and suboptimal. While suboptimal herding is the greatest concern from a financial stability perspective, optimal herding can build up financial vulnerabilities as well.





Figure 5: Simulated pricing errors: AI and human behaviors

The figure shows the average (left) and median (right) absolute pricing errors across trading periods simulated in the experimental model under (a) Treatment I and (b) Treatment II with 1000 independent iterations. The lines show four different behavioral assumptions: trading decisions following the distributions uncovered from the experiments on AI agents, humans, an equal mix between AI agents and humans, and Optimal AI agents. Details on the simulations are provided in Appendix B.

(a) Treatment I



(b) Treatment II




Table 1: Trading behavior in AI and human laboratories

The table shows the distribution of decisions in the human and AI laboratories. Decisions are averaged across all sessions and trading periods in (a) Treatment I (without event uncertainty), (b) Treatment II (with event uncertainty), and (c) Treatment III (without price updating). “Human” decisions are taken directly from Cipriani and Guarino (2009) for Treatment I and II. “AI” decisions represent the average decisions across all LLMs. The table also show the results separately for each LLM. “Rational” behavior represents cases where the informed trader chooses to buy upon receiving a white signal and sell upon receiving a blue signal. “Partial Rational” behavior represents cases where the informed trader chooses to buy (sell) upon receiving a white (blue) signal and not trade upon receiving the other signal. “Cascade Trading” represents cases where the informed trader chooses the same trading action (buy or sell) regardless of the private signal. These decisions are decomposed into “Optimal Herding”, “Suboptimal Herding”, “Contrarian” behavior, and cases where the trade imbalance is zero (“Undetermined”). While the exact decomposition of human cascade trading decisions is unknown, we simulate the decomposition given the distributions of trading decisions from Cipriani and Guarino (2009) using 10,000 independent iterations. “Cascade No Trading” represents cases where the informed trader chooses not to trade regardless of the private signal. “Error” represents cases where the informed trader chooses to buy upon receiving a blue signal and sell upon receiving a white signal. The table also reports the frequency of trading periods where herding is optimal.

(a) Treatment I

	Human	AI	Claude 3.7	Claude 3.5	Llama 3	Nova Pro
Rational	45.70%	61.00%	70.97%	37.11%	97.66%	38.28%
Partial Rational	19.60%	29.48%	11.29%	45.31%	0.00%	61.33%
Cascade Trading	19.00%	9.42%	17.74%	17.58%	2.34%	0.00%
Optimal Herding	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Suboptimal Herding	4.85%	0.00%	0.00%	0.00%	0.00%	0.00%
Contrarian	8.90%	9.42%	17.74%	17.58%	2.34%	0.00%
Undetermined	5.20%	0.00%	0.00%	0.00%	0.00%	0.00%
Cascade No Trading	12.30%	0.10%	0.00%	0.00%	0.00%	0.39%
Error	3.40%	0.00%	0.00%	0.00%	0.00%	0.00%
Optimal Herding Opportunities	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

(b) Treatment II

	Human	AI	Claude 3.7	Claude 3.5	Llama 3	Nova Pro
Rational	50.90%	97.36%	100.00%	100.00%	100.00%	89.45%
Partial Rational	20.10%	2.64%	0.00%	0.00%	0.00%	10.55%
Cascade Trading	12.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Optimal Herding	4.12%	0.00%	0.00%	0.00%	0.00%	0.00%
Suboptimal Herding	2.08%	0.00%	0.00%	0.00%	0.00%	0.00%
Contrarian	2.41%	0.00%	0.00%	0.00%	0.00%	0.00%
Undetermined	3.37%	0.00%	0.00%	0.00%	0.00%	0.00%
Cascade No Trading	16.50%	0.00%	0.00%	0.00%	0.00%	0.00%
Error	0.50%	0.00%	0.00%	0.00%	0.00%	0.00%
Optimal Herding Opportunities	39.50%	36.56%	30.61%	46.88%	21.88%	46.88%

(c) Treatment III

	AI	Claude 3.7	Claude 3.5	Llama 3	Nova Pro
Rational	99.65%	99.38%	99.22%	100.00%	100.00%
Partial Rational	0.16%	0.62%	0.00%	0.00%	0.00%
Cascade Trading	0.00%	0.00%	0.00%	0.00%	0.00%
Optimal Herding	0.00%	0.00%	0.00%	0.00%	0.00%
Suboptimal Herding	0.00%	0.00%	0.00%	0.00%	0.00%
Contrarian	0.00%	0.00%	0.00%	0.00%	0.00%
Undetermined	0.00%	0.00%	0.00%	0.00%	0.00%
Cascade No Trading	0.20%	0.00%	0.78%	0.00%	0.00%
Error	0.00%	0.00%	0.00%	0.00%	0.00%
Optimal Herding Opportunities	38.36%	50.31%	34.38%	34.38%	34.38%

Table 2: LLM analysis of AI agent reasoning

The table shows LLM analysis of reasoning passages using the Claude 3.7 model. The model is prompted to read each passage of reasoning and answer the following five questions. Question 1: Is the trader comparing the price to the expected fundamental value of the asset? (True/False). Question 2: Is the expected value computed using only the signal accuracy and the signal, e.g., $0.7*100+0*0.3=70$ or $0.7*0+0.3*100=30$? (True/False). Question 3: Does the trader consider the market trend or the trading history in their reasoning? (True/False). Question 4: How does the trader characterize the attractiveness of the investment? Question 5: On a scale from 0-100 (where 100 represents purely emotional and 0 represents purely rational or logical), how much is the investor driven by emotions in their assessment? For “True/False” questions, the table reports the fraction of “True” answers. Panel (a) reports the analysis of reasoning provided by all LLMs and panel (b) reports analysis of reasoning in the Llama 3 model.

(a) AI

	Treatment I	Treatment II	Treatment III
Question 1	99.16%	99.01%	99.67%
Question 2	63.07%	63.09%	63.51%
Question 3	17.05%	9.50%	24.12%
Question 4			
VERY ATTRACTIVE	11.69%	1.88%	0.71%
ATTRACTIVE	42.83%	69.39%	68.33%
REASONABLE	12.65%	6.08%	7.33%
LESS ATTRACTABLE	10.36%	3.65%	5.69%
NO INCENTIVE	22.17%	19.01%	17.89%
Question 5			
Mean	6.39%	4.93%	5.62%
Bottom decile	0.00%	0.00%	0.00%
Median	0.00%	0.00%	0.00%
Top decile	20.00%	15.00%	20.00%

(b) Llama 3

	Treatment I	Treatment II	Treatment III
Question 1	97.46%	100.00%	99.61%
Question 2	13.87%	4.49%	4.30%
Question 3	40.62%	30.66%	66.60%
Question 4			
VERY ATTRACTIVE	1.56%	1.56%	0.39%
ATTRACTIVE	69.92%	83.79%	74.61%
REASONABLE	13.48%	4.88%	4.49%
LESS ATTRACTABLE	10.16%	5.66%	8.98%
NO INCENTIVE	4.69%	4.10%	11.52%
Question 5			
Mean	14.53%	12.72%	16.65%
Bottom decile	5.00%	5.00%	10.00%
Median	15.00%	10.00%	15.00%
Top decile	25.00%	20.00%	25.00%

Table 3: Trading behavior in optimal AI laboratory

The table shows the distribution of decisions in the optimal AI laboratory in which LLMs are prompted with guidance on optimal decision making. Decisions are averaged across all sessions and trading periods in (a) Treatment I (without event uncertainty), (b) Treatment II (with event uncertainty), and (c) Treatment III (without price updating). “Optimal AI” decisions represent the average decisions across all LLMs. The table also show the results separately for each LLM. “Rational” behavior represents cases where the informed trader chooses to buy upon receiving a white signal and sell upon receiving a blue signal. “Partial Rational” behavior represents cases where the informed trader chooses to buy (sell) upon receiving a white (blue) signal and not trade upon receiving the other signal. “Cascade Trading” represents cases where the informed trader chooses the same trading action (buy or sell) regardless of the private signal. These decisions are decomposed into “Optimal Herding”, “Suboptimal Herding”, “Contrarian” behavior, and cases where the trade imbalance is zero (“Undetermined”). “Cascade No Trading” represents cases where the informed trader chooses not to trade regardless of the private signal. “Error” represents cases where the informed trader chooses to buy upon receiving a blue signal and sell upon receiving a white signal. The table also reports the frequency of trading periods where herding is optimal.

(a) Treatment I

	Optimal AI	Claude 3.7	Claude 3.5	Llama 3	Nova Pro
Rational	55.88%	48.15%	52.34%	95.31%	27.73%
Partial Rational	40.60%	51.85%	38.28%	0.00%	72.27%
Cascade Trading	3.52%	0.00%	9.38%	4.69%	0.00%
Optimal Herding	0.00%	0.00%	0.00%	0.00%	0.00%
Suboptimal Herding	0.00%	0.00%	0.00%	0.00%	0.00%
Contrarian	3.52%	0.00%	9.38%	4.69%	0.00%
Undetermined	0.00%	0.00%	0.00%	0.00%	0.00%
Cascade No Trading	0.00%	0.00%	0.00%	0.00%	0.00%
Error	0.00%	0.00%	0.00%	0.00%	0.00%
Optimal Herding Opportunities	0.00%	0.00%	0.00%	0.00%	0.00%

(b) Treatment II

	Optimal AI	Claude 3.7	Claude 3.5	Llama 3	Nova Pro
Rational	18.65%	37.09%	12.50%	12.50%	12.50%
Partial Rational	21.88%	0.00%	0.00%	0.00%	87.50%
Cascade Trading	59.48%	62.91%	87.50%	87.50%	0.00%
Optimal Herding	47.43%	39.74%	75.00%	75.00%	0.00%
Suboptimal Herding	0.00%	0.00%	0.00%	0.00%	0.00%
Contrarian	6.60%	13.91%	6.25%	6.25%	0.00%
Undetermined	5.44%	9.27%	6.25%	6.25%	0.00%
Cascade No Trading	0.00%	0.00%	0.00%	0.00%	0.00%
Error	0.00%	0.00%	0.00%	0.00%	0.00%
Optimal Herding Opportunities	81.52%	63.58%	87.50%	87.50%	87.50%

(c) Treatment III

	Optimal AI	Claude 3.7	Claude 3.5	Llama 3	Nova Pro
Rational	51.05%	45.60%	31.25%	35.55%	91.80%
Partial Rational	0.00%	0.00%	0.00%	0.00%	0.00%
Cascade Trading	48.95%	54.40%	68.75%	64.45%	8.20%
Optimal Herding	44.36%	53.60%	59.38%	56.25%	8.20%
Suboptimal Herding	4.01%	0.80%	7.03%	8.20%	0.00%
Contrarian	0.00%	0.00%	0.00%	0.00%	0.00%
Undetermined	0.59%	0.00%	2.34%	0.00%	0.00%
Cascade No Trading	0.00%	0.00%	0.00%	0.00%	0.00%
Error	0.00%	0.00%	0.00%	0.00%	0.00%
Optimal Herding Opportunities	50.90%	53.60%	59.38%	56.25%	34.38%

Table 4: Expected Payoffs

The table shows descriptive statistics of expected payoffs in the AI and Optimal AI laboratories. In the Optimal AI laboratory, LLMs are prompted with guidance on optimal decision making. Expected payoffs are computed as $\mathbb{E}(v|s_t, h_t) - p_t$ if the agent decides to buy the asset, $p_t - \mathbb{E}(v|s_t, h_t)$ if the agent decides to sell the asset, and zero otherwise. Expected payoffs are averaged across all sessions, trading periods, and LLMs before computing the statistics.

	Treatment I		Treatment II		Treatment III	
	AI	Optimal AI	AI	Optimal AI	AI	Optimal AI
Mean	2.57	2.72	3.80	14.95	5.07	7.79
Median	2.74	2.74	6.67	19.53	6.67	11.49
Min	-6.67	-6.67	-11.44	-28.28	-16.19	-16.19
Max	6.67	6.67	11.55	28.35	16.46	16.63
Std Dev	3.90	3.57	6.47	14.20	8.83	7.87

Table 5: Trading behavior in AI laboratory with different types of signals

The table shows the distribution of decisions in the AI laboratory when varying the color used to code the private information signals. Respectively the white and blue signals are replaced by (i) green and red, and (ii) red and green. Decisions are averaged across all sessions, trading periods, and LLMs in (a) Treatment I (without event uncertainty), (b) Treatment II (with event uncertainty), and (c) Treatment III (without price updating). “Rational” behavior represents cases where the informed trader chooses to buy upon receiving a white signal and sell upon receiving a blue signal. “Partial Rational” behavior represents cases where the informed trader chooses to buy (sell) upon receiving a white (blue) signal and not trade upon receiving the other signal. “Cascade Trading” represents cases where the informed trader chooses the same trading action (buy or sell) regardless of the private signal. These decisions are decomposed into “Optimal Herding”, “Suboptimal Herding”, “Contrarian” behavior, and cases where the trade imbalance is zero (“Undetermined”). “Cascade No Trading” represents cases where the informed trader chooses not to trade regardless of the private signal. “Error” represents cases where the informed trader chooses to buy upon receiving a blue signal and sell upon receiving a white signal. The table also reports the frequency of trading periods where herding is optimal.

(a) Treatment I

	Good: Green, Bad: Red	Good: Red, Bad: Green
Rational	54.83%	20.23%
Partial Rational	37.55%	42.76%
Cascade Trading	7.62%	6.64%
Optimal Herding	0.00%	0.00%
Suboptimal Herding	0.00%	3.52%
Contrarian	7.62%	0.00%
Undetermined	0.00%	3.12%
Cascade No Trading	0.00%	5.18%
Error	0.00%	25.20%
Optimal Herding Opportunities	0.00%	0.00%

(b) Treatment II

	Good: Green, Bad: Red	Good: Red, Bad: Green
Rational	98.54%	50.78%
Partial Rational	1.46%	11.72%
Cascade Trading	0.00%	12.50%
Optimal Herding	0.00%	7.32%
Suboptimal Herding	0.00%	1.56%
Contrarian	0.00%	0.00%
Undetermined	0.00%	3.61%
Cascade No Trading	0.00%	0.00%
Error	0.00%	25.00%
Optimal Herding Opportunities	52.94%	42.93%

(c) Treatment III

	Good: Green, Bad: Red	Good: Red, Bad: Green
Rational	99.64%	23.24%
Partial Rational	0.00%	25.00%
Cascade Trading	0.36%	25.18%
Optimal Herding	0.36%	18.82%
Suboptimal Herding	0.00%	3.14%
Contrarian	0.00%	0.18%
Undetermined	0.00%	3.04%
Cascade No Trading	0.00%	0.00%
Error	0.00%	26.58%
Optimal Herding Opportunities	56.61%	56.75%

Appendices

A. Prompts

Prompt 1: System prompt

This prompt describes the instructions of the experiment, which is given to the LLMs through their system prompt.

You are participating in an experiment at the Experimental Laboratory of the ELSE Centre at the
↪ Department of Economics at UCL. The instructions given for the laboratory experiment are as
↪ follows:

There are a total of 8 participants in this experiment. Everyone is receiving the same
↪ instructions.

In the experiment, you can exchange one unit of an asset with a computerized market maker. You
↪ and the other participants will make trading decisions through 8 sequential rounds. In each
↪ round, only one participant will be selected to trade. Each participant can only trade
↪ once.

In each round, the market maker sets the price of the asset as the expected value of the
↪ fundamental value of the asset, conditional on the history of the trades from the previous
↪ rounds.

[if treatment==2: {The market maker will update the price as if, with high probability, it were
↪ trading not with informed traders, but with noise traders.}]

The fundamental value of the asset is a discrete random variable that can take values 0 or 100,
↪ each with a 50% probability. You do not know the fundamental value of the asset, but you
↪ may receive a signal (white or blue) on the value. If the asset value is 100, you receive a
↪ white signal with 70% probability and a blue signal with 30% probability. If the value is
↪ 0, you receive a white signal with 30% probability and a blue signal with 70% probability.

You will be making decisions on whether to buy or sell one unit of the asset at a given price,
↪ or not to trade given respectively a white and a blue signal. The realized signal will only
↪ be revealed to you if you are selected to trade. After each round, the computer will
↪ randomly select a participant whose trade gets executed. That participant receives the
↪ realized signal. The remaining participants then observe the executed trading decision
↪ (buy, sell, or no trade), but do not receive the realized signal. They also do not observe
↪ the identity of the selected participant. The procedure continues for 8 rounds until all
↪ participants have acted once. All participants (including those whose decision has already
↪ been executed) observe the trading decisions in each period and the corresponding price
↪ movement.

After 8 rounds, the asset value is revealed, and each participant receives a payoff computed
↪ based on the trading decision and price in the round in which the participant was selected
↪ and the asset value v .

Payoffs are computed in a fictitious experimental currency called lira. If the participant sold
↪ the asset at price p , the payoff is $p-v$ lira. If the participant bought the asset at price
↪ p , the payoff is $v-p$ lira. If the participant decided not to trade, the payoff is zero
↪ lira. At the end of the experiment, the payoffs are added up and converted into British
↪ pounds at the rate of 3 lira per pound. In addition, you are paid 70 pounds for
↪ participating in the experiment, regardless of your payoff.

Prompt 2: User prompt in AI laboratory

This prompt describes the instructions given to the LLMs in each trading period to each agent j . The HISTORY input consists of the executed trades of selected traders along with the history of actions and reasoning for agent j in all previous periods. In addition to this user prompt, the LLMs have available the instructions through the system prompt, see Prompt 1.

```
This is round [TRADING PERIOD (t)].
```

```
[HISTORY]
```

```
If you receive a white signal, will you buy, sell, or not trade at a price of [PRICE]?
```

```
If you receive a blue signal, will you buy, sell, or not trade at a price of [PRICE]?
```

```
Please make sure that you provide your response in the following format:
```

```
{  
  "actionWhite": "BUY/SELL/NO TRADE at the price of [PRICE] conditional on observing a white  
    ↪ signal",  
  "actionBlue": "BUY/SELL/NO TRADE at the price of [PRICE] conditional on observing a blue  
    ↪ signal",  
  "reasoningWhite": "Brief explanation of your decision conditional on observing a white  
    ↪ signal (1-2 sentences) ",  
  "reasoningBlue": "Brief explanation of your decision conditional on observing a blue signal  
    ↪ (1-2 sentences)"  
}
```

Prompt 3: User prompt in optimal AI laboratory

This prompt describes the instructions given to the LLMs in each trading period to each agent j in the optimal AI laboratory. The HISTORY input consists of the executed trades of selected traders along with the history of actions and reasoning for agent j in all previous periods. In addition to this user prompt, the LLMs have available the instructions through the system prompt, see Prompt 1.

```
This is round [TRADING PERIOD (t)].

if TRADING PERIOD (t)==1:
    Note that given current conditions, it is optimal to buy given a white signal and sell
    ↪ given a blue signal.

else:
    [HISTORY]

    if expected_value_trader_white > price and expected_value_trader_blue > price:
        Note that given current conditions, it is optimal to follow the herd and buy regardless
        ↪ of the signal.

    if expected_value_trader_white < price and expected_value_trader_blue < price::
        Note that given current conditions, it is optimal to follow the herd and sell
        ↪ regardless of the signal.

If you receive a white signal, will you buy, sell, or not trade at a price of [PRICE]?
If you receive a blue signal, will you buy, sell, or not trade at a price of [PRICE]?

Please make sure that you provide your response in the following format:
{
    "actionWhite": "BUY/SELL/NO TRADE at the price of [PRICE] conditional on observing a white
    ↪ signal",
    "actionBlue": "BUY/SELL/NO TRADE at the price of [PRICE] conditional on observing a blue
    ↪ signal",
    "reasoningWhite": "Brief explanation of your decision conditional on observing a white
    ↪ signal (1-2 sentences) ",
    "reasoningBlue": "Brief explanation of your decision conditional on observing a blue signal
    ↪ (1-2 sentences)"
}
```

Prompt 4: System prompt personal characteristics add-on

This prompt describes an add-on to the system prompt that provides characteristics of the AI agent. The characteristics are drawn randomly from the unconditional distributions of human participant characteristics reported in Cipriani and Guarino (2009) restricted according to a set of heuristics to ensure realistic personas.

```
You are a [AGE]-year old [GENDER]. You work as a [OCCUPATION] and you have [TENURE] years of
↪ tenure. You have a [EDUCATION LEVEL] degree in [EDUCATION FIELD]. Respond in way that is
↪ consistent with the knowledge and expected behavior of a person with these characteristics.
```

B. Details on price convergence simulations

We simulate pricing dynamics over $T = 100$ trading periods for type each agent (AI, human, Optimal AI) separately, and for a mixture of human and AI agents. We fix the distribution of trading decisions for each type as reported in Table 1. We simulate 1000 independent iterations. For each iteration, the simulation performs the following steps:

1. Assume that an information event has occurred²¹ and draw an asset value $v \in \{0, 100\}$ with equal probability.
2. Set the probabilities of each signal based on the sampled asset value. If $v = 0$, $\Pr(\text{white signal}) = 0.3$ and $\Pr(\text{blue signal}) = 0.7$. If $v = 100$, $\Pr(\text{white signal}) = 0.7$ and $\Pr(\text{blue signal}) = 0.3$.
3. Set the initial price at 50.
4. For each trading round, $t = 1, 2, \dots, T$, where $T = 100$:
 - (a) For each participant, draw a trading decision from the distribution of decisions:
 - Determine if herding is optimal in the current round given the price and trading history.
 - If herding is optimal, sample a random variable determining if the participant chooses to herd. The probability of herding in this case is given by the ratio of the percent of decisions where the trader engages in optimal herding to the percentage of decisions where herding is optimal, both available from Table 1.
 - If the participant does not herd, sample a decision from the set of remaining decisions, {Rational, Partially Rational, Contrarian, Cascade No Trading, Error}. Probabilities of each type of decision is given in Table 1. The probabilities are normalized to represent a valid probability distribution whose entries sum to one.
 - If herding is not optimal, draw a random decision from the set of remaining decisions, {Rational, Partially Rational, Suboptimal Herding, Contrarian, Undetermined, Cascade No Trading, Error}. Probabilities of each type of decision is given in Table 1. The probabilities are normalized to represent a valid probability distribution whose entries sum to one.
 - Based on the sampled type of decision, set the participant's trading decision ($d \in \{\text{buy}, \text{sell}, \text{notrade}\}$) as follows:

21 This is consistent with how we and Cipriani and Guarino (2009) implement the experiments: even in Treatment II, where the event probability is 15%, the experiment is conducted under the assumption that an event did occur.

- *Rational*: Sample $d \in \{\text{buy}, \text{sell}\}$ with probabilities $\Pr(\text{white signal})$ and $\Pr(\text{blue signal})$.
- *Partially Rational*: With equal probability, either sample $d \in \{\text{buy}, \text{no trade}\}$ with probabilities $\Pr(\text{white signal})$ and $\Pr(\text{blue signal})$ or sample $d \in \{\text{no trade}, \text{sell}\}$ with probabilities $\Pr(\text{white signal})$ and $\Pr(\text{blue signal})$.
- *Optimal Herding*: Set d equal to the mode of the trading history.
- *Suboptimal Herding*: Set d equal to the mode of the trading history.
- *Contrarian*: Set d equal to the opposite trading action than the mode of the trading history.
- *Undetermined*: Sample $d \in \{\text{no trade}, \text{sell}\}$ with equal probabilities.
- *Cascade No Trading*: Set $d = \text{no trade}$.
- *Error*: Sample $d \in \{\text{buy}, \text{sell}\}$ with probabilities $\Pr(\text{blue signal})$ and $\Pr(\text{white signal})$.

(b) Randomly select a trader (without replacement) to execute their decision.

(c) Based on the executed trade, update the price and the trading history.

Repeating these steps 1000 times yield a sample of prices throughout 100 trading periods given known asset values from which we can compute pricing errors. The average absolute pricing errors for each trading period are shown in Figure 5 in the main text. The pricing errors converge to zero with the exception of the case where traders act according to the Optimal AI decision probabilities. Table A1 provides more insights by showing the distributions of the 100th period pricing errors across all iterations. These results highlight that Optimal AI pricing errors do not converge to zero on average due to a few outlier cases, where the price converge to either zero with a value of 100, or vice versa. These cases arise because, as in the real world, signals about fundamentals are not fully accurate.

Table A1: Pricing error quantiles at the 100th trading period

The table shows quantiles of the absolute pricing errors at the 100th trading period simulated in the experimental model under (a) Treatment I and (b) Treatment II with 1000 independent iterations.

(a) Treatment I

	0%	25%	50%	75%	100%
AI	0.00	0.00	0.00	0.00	99.38
Human	0.00	0.00	0.00	0.00	99.89
Human-AI	0.00	0.00	0.00	0.00	84.48
Optimal AI	0.00	0.00	0.00	0.00	7.30

(b) Treatment II

	0%	25%	50%	75%	100%
AI	0.01	0.05	0.19	0.51	4.57
Human	0.02	0.17	0.37	0.82	78.24
Human-AI	0.01	0.08	0.19	0.59	34.53
Optimal AI	0.00	0.01	0.04	0.13	99.99