Does Your AI Move Your Money?

Adoption and Impact of a Generative AI Investment Assistant

Longjiang He, Zhenhua Li, Jinfei Sheng, Ying Yu, and Xiaoyan Zhang*

August 2026

Abstract

We study how mutual fund investors interact with a generative-AI investment assistant (GenAI Assistant) on a large Chinese fintech platform. Using account-level data with detailed records of investor queries and AI-generated responses, we provide the first large-scale field evidence on AI financial advice in deployment: what investors ask, what the AI recommends and why, whether the advice moves trades, and whether it has investment value. We find that investors underserved by traditional advisory channels, especially those with less education and lower income, are more likely to adopt GenAI Assistant. The assistant's recommendations are tailored: they respond to investor risk profiles and screen strongly against high-fee funds, though they also chase past returns. Importantly, AI advice moves behavior: investors follow the AI's recommendations in 51% of interactions within 24 hours, and they are more likely to follow the advice when the recommendation is stronger and the response is more relevant and readable. Most strikingly, GenAI advice has investment value through trading discipline, not fund selection. Recommended funds do not outperform, but investors who follow the advice trade far less and earn higher returns and Sharpe ratios at no additional risk. Overall, our results speak to the broader debate over GenAI's economic value: in finance, its deployment delivers measurable benefits to retail investors.

Keywords: Financial AI, mutual fund investors, big data, LLMs, retail investors, AI assistant, companionship, trading behavior.
JEL Codes: C52, C5, G1, G14.

---

* Longjiang He (helj.22@pbcsf.tsinghua.edu.cn), Ying Yu (yuy.22@pbcsf.tsinghua.edu.cn) and Xiaoyan Zhang (zhangxiaoyan@pbcsf.tsinghua.edu.cn) are at the PBC School of Finance at Tsinghua University. Zhenhua Li (sunny.lzh@alibaba-inc.com) is from Ant Group. Jinfei Sheng (jinfei.sheng@uci.edu) is at Merage School of Business, University of California, Irvine. Xiaoyan Zhang gratefully acknowledges financial support from the National Natural Science Foundation of China [Grant No. 72350710220]. All remaining errors are our own. Corresponding author: Xiaoyan Zhang, PBC School of Finance, 43 Chengfu Road, Beijing, China, 100083, zhangxiaoyan@pbcsf.tsinghua.edu.cn.

* The authors acknowledge and appreciate the supports from the Digital Economy Open Research Platform (www.deor.org.cn). All data was sampled and desensitized, and was remotely analyzed on the Ant Open Research Laboratory in an Ant Group Environment which is only remotely accessible for empirical analysis.

1. Introduction

Generative artificial intelligence (GenAI) has attracted extraordinary attention and capital. Since the launch of large language models (LLMs), investment has poured into computing infrastructure, including chips, data centers, and model training, on the premise that the technology will transform the economy, and equity markets have priced this promise into trillions of dollars of valuation. Yet an intense debate has emerged over whether the promise will be realized: skeptics note that measurable productivity gains remain scarce, that adoption inside firms is shallower than headline figures suggest, and that revenues from AI applications lag far behind the investment in AI capability. At the center of this debate is a simple question that has received surprisingly little direct evidence: where are the use cases? For all that is known about what GenAI can do in benchmarks and demonstrations, little is known about what it does do when deployed to real users making real decisions: who uses it, how they use it, whether they act on its output, and whether doing so makes them better off. Motivated by this debate, we examine a use case of GenAI in finance.

Recent advances in GenAI are rapidly reshaping the financial services industry and are increasingly embedded in retail investment platforms (Goldstein, Jiang, and Karolyi, 2019). Unlike traditional robo-advisors, which primarily rely on rule-based portfolio allocation and standardized recommendations, generative AI systems provide real-time, customized, interactive responses that integrate information, investment advice, and emotional support. By tailoring responses to user queries, profiles, and market conditions, these systems act as both information and advice intermediaries in the investment process. These features are particularly relevant in the context of retail investing, where individuals face several well-documented frictions, such as information collection, analytical sophistication, and cognitive constraints, (Gomes, Haliassos, and Ramadorai, 2021), and where even routine portfolio decisions, such as rebalancing and asset allocation, involve substantial complexity (Dammon, Spatt, and Zhang, 2004). Yet as AI-fueled financial services advance, many important questions remain unanswered: Who uses AI advisory services? What do these assistants actually recommend,

and is the advice any good? Do investors follow the recommendations and does following them improve investment performance?

The main challenge for answering these questions is data. Because AI services are still early in development, providers rarely make interaction data available to researchers, so existing evidence on what GenAI *can* do in benchmarks and experiments far outstrips evidence on what it *does* do in the hands of real users. Large-scale data from financial platforms have become a central engine of new research in finance (Goldstein, Spatt, and Ye, 2021), and our setting extends this frontier to AI-investor interactions. We close this gap through access to the largest Chinese fintech platform. We study how mutual fund investors interact with a generative-AI investment assistant deployed at scale on this platform, observing, for the first time in the field, the complete chain from what investors ask, to what the AI recommends, to whether the advice moves trades, to whether it creates investment value. Specifically, we collect account-level data on 400,000 randomly selected active mutual fund investors from July 2024 to October 2025, with detailed records of investor characteristics, AI queries, AI-generated responses, subsequent trading, and investment performance.

We first examine who adopts the AI investment assistant by asking the AI assistant questions. As a low-cost source of financial guidance, AI can provide the financial literacy that sound investment decisions otherwise require. Interestingly, the financial literacy literature documents that less literate households make costlier investment mistakes and are less likely to obtain professional advice (Lusardi and Mitchell, 2014; van Rooij, Lusardi, and Alessie, 2011). If AI advice lowers the cost of guidance, it should appeal most to precisely these disadvantaged investors. Using investor-level regressions of adoption and usage intensity on investor characteristics, we find strong support for this prediction: controlling for a rich set of investor characteristics, adoption is significantly higher among investors with lower education and lower household income. The magnitudes are economically meaningful relative to the 31% baseline adoption rate: investors without a bachelor's degree are 2.8 percentage points more likely to adopt and submit 7% more queries, and adoption declines steadily with household income, with investors in the lowest income category being 6 percentage points more likely to

adopt than those in the highest. The same patterns holding for both information-seeking and advice-seeking queries. This adoption profile inverts the gradient documented for robo-advisors and human advisors, whose users tend to be wealthier and more sophisticated; to our knowledge, ours is the first evidence that an advice technology organically attracts the low-literacy investors whom the advice gap literature identifies as needing guidance most. Who shares in the gains from algorithmic financial technology is a first-order question (Fuster, Goldsmith-Pinkham, Ramadorai, and Walther, 2022), and our adoption evidence speaks directly to it.

We next examine what the AI recommends and whether its advice is customized, estimating how the direction and strength of recommendations vary with investor and fund characteristics. We find that AI responses follow a clear structured communication pattern: first emotional support, next information provision, then recommendation and conclusion. In addition, recommendations are systematically tailored on both sides of the interaction. On the investor side, the assistant gives more aggressive advice to risk-tolerant investors and more conservative advice to wealthier investors and those holding riskier, equity-heavy portfolios, conditioning on suitability in the spirit of fiduciary advice. On the fund side, the assistant's strongest tilt is against fees: it strongly disfavors expensive funds, and it also penalizes volatile funds while chasing past returns. To our knowledge, ours is the first evidence on the content of AI advice at scale, revealing an assistant that conditions on the individual investor and tilts against fund costs, though its return-chasing previews the predictability results below.

We then ask whether investors follow the advice, and what makes the advice more likely to be followed. Our hypothesis is that if investors regard the AI as a credible advisor, they would trade in the direction it recommends; and advice would be followed with a higher probability if more strongly justified (Asquith, Mikhail and Au, 2005) and accessible and clearly directional under limited attention (Peng and Xiong, 2006). Linking each recommendation to the investor's trades within 24 hours, we find strong support for both predictions. Investors follow the AI's recommendations in 51% of interactions, and trading tilts significantly toward the recommended direction: moving from a hold to a strong buy

recommendation raises the net propensity to buy by 2.4 percentage points, economically large against the roughly 14% unconditional probability of trading after a query. The choice of "follow AI recommendation" responds to the message itself: a one-standard-deviation increase in recommendation strength raises following by 1.3 percentage points, with smaller but significant effects for relevance and readability. The choice of follow is also more common among less-educated and explicitly advice-seeking investors. This provides the missing individual-level link in the emerging evidence on AI and trading: existing studies infer reliance from aggregates, such as trading declines during ChatGPT outages (Cheng, Lin and Zhao, 2025) or abnormal volume around AI query spikes (Ecker, Li, Li and Wu, 2026). By connecting specific recommendations to specific trades by identified investors, we show directly that AI advice moves individual behavior and, echoing the analyst-recommendation literature's findings for professional advice, that its persuasiveness depends on justification and presentation, now measured at the level of individual response characteristics.

Finally, we examine whether the advice creates investment value, testing two competing predictions. If AI synthesizes information beyond what prices incorporate, recommended funds should outperform, which is consistent with evidence that LLMs can forecast returns (Chen, Kelly, and Xiu, 2022; Chen, Tang, Zhou, and Zhu, 2025; He, Tan, and Zhang, 2026). Alternatively, recommendations distilled from public information should contain no alpha, yet the advice could still create value by improving investor behavior, given retail investors' well-documented costly overtrading (Barber and Odean, 2000). Our evidence decisively favors the behavioral view. AI recommendations do not predict fund returns at any horizon, at the query level or aggregated to fund-day ratios, with economically small point estimates. Yet, investors who follow the advice earn portfolio returns 11 basis points higher over the following month, roughly 1.3% annualized, with higher Sharpe ratios and no increase in volatility. Decomposition by recommendation type reveals the mechanism: gains are concentrated entirely among investors who follow hold recommendations, whose next-month returns are 17.5 basis points higher, while following buy recommendations yields no gains. The channel is turnover, which falls by 13.6 percentage points of portfolio value per month for followers,

implying transaction-cost savings of the same order as the return improvement. This mechanism is new to the digital-advice literature. Robo-advisors improve outcomes primarily through diversification and rebalancing of delegated portfolios (D'Acunto, Prabhala, and Rossi, 2019; Rossi and Utkus, 2024); LLM-based advice has been shown to raise returns among adherent users in experimental settings (Lu, Li, and Shi, 2025). We show that in organic deployment, where the AI executes nothing and investors choose freely, its value operates through a distinct channel: persuading investors not to trade. The AI creates value not by telling investors what to buy, but by disciplining when they act at all.

Our study contributes to three strands of literature. First, we contribute to the rapidly emerging literature on the adoption and economic consequences of generative AI. Recent studies document the rapid growth of LLM usage across the economy and its increasing role in financial decision-making (Chatterji, Cunningham and Deming, 2025; Bick, Blandin and Deming, 2026; Ecker, Li, Li and Wu, 2026). Within finance, a growing body of work examines what generative AI *can* do: building on earlier evidence that machine learning methods can uncover return predictability (Chinco, Clark-Joseph, and Ye, 2019), LLMs can forecast stock price movements from public news (Lopez-Lira and Tang, 2026) and predict returns across markets (Chen, Kelly, and Xiu, 2022; He, Tan, and Zhang, 2026), and asset managers that adopt generative AI improve fund performance (Sheng, Sun, Yang, and Zhang, 2026). A parallel literature examines AI as a market participant, including its potential for algorithmic collusion, market fragility, and competition with human traders (Dou, Goldstein, and Ji, 2025, 2026; Dou, Goldstein, Li, and Yang, 2026); we study AI as an advisor to households. Evidence that investors rely on these tools is accumulating but largely indirect: Cheng, Lin and Zhao (2025) show that stock trading volume declines during ChatGPT outages, and Ecker, Li, Li and Wu (2026) link AI query volume to abnormal trading activity. By connecting specific recommendations to specific trades by identified investors, we provide the missing individual-level evidence that AI advice moves behavior, and that its persuasiveness depends on justification and presentation. Our results also add a cautionary counterpoint to the capability literature: while LLMs can forecast returns in research settings, the recommendations of a

deployed consumer-facing assistant contain no predictive power for fund returns — the value of its advice arises through investor behavior instead.

Second, this paper is linked to studies on digital advice and robo-advisors. Existing studies (e.g., D'Acunto, Prabhala, and Rossi, 2019; Rossi and Utkus, 2024; Bianchi and Brière, 2026) primarily examine automated portfolio allocation tools grounded in classical portfolio optimization, which improve outcomes chiefly through diversification and rebalancing of delegated portfolios, and whose adopters tend to be relatively sophisticated. Generative-AI-based advisory systems differ on both margins. They interact with investors through natural-language conversation by providing information, explanations, recommendations, and emotional support rather than constructing or executing portfolios, so any effect must operate through persuasion rather than delegation. Their adopters differ correspondingly: usage concentrates among less-educated, lower-income investors, inverting the sophistication gradient documented for robo-advice. We contribute to this literature by providing the first large-scale evidence on how investors use conversational financial AI and by identifying a distinct value channel: where robo-advisors improve portfolios by taking control of them, generative AI improves outcomes by persuading investors — most importantly, persuading them not to trade.

Closest to our work, Lu, Li, and Shi (2025) conduct an opt-in randomized experiment in which an LLM explains the recommendations of a pre-existing, rule-based robo-advisor to roughly 29,000 self-selected investors at one brokerage, and show that the LLM's explanations improve adherence and outcomes. Unlike their intervention study, ours is a deployment study. The AI's role differs fundamentally across the two papers: in their setting, the LLM explains advice generated elsewhere, whereas in ours the AI is the advisor, generating personalized fund-level recommendations for millions of ordinary investors with no experimental framing and no prescribed portfolio behind it. The two papers accordingly answer complementary questions. Their randomization cleanly establishes that LLM-enhanced advice can improve outcomes for investors who sign up for it; our setting reveals what AI advice actually is and does in the reality: who adopts it when it is simply available, what the AI recommends and

why, how persuasively it moves trades, and where its value comes from. These questions require observing the full query-response-trade chain at the population level, which an opt-in experiment cannot produce.

2. Data

2.1 The AI Investment Assistant

Our data come from *Ant Fortune*, the largest online wealth management platform in China. By the end of 2025, *Ant Fortune* serves more than 500 million investors[1], with RMB 1.81 trillion in non-money-market mutual fund holdings. *Ant Fortune* provides mutual fund products to retail investors, with more than 10,000 mutual funds on its platform, covering more than 90 % of the total market cap of China's non-money-market mutual fund industry.

Established in year 2015, *Ant Fortune* has been one of the pioneers of using technology to provide financial services, with a mission statement of "making professional and personalized wealth management services more accessible to individual investors." In 2019, the platform introduced an automatic robo-advisor, *Zhi Xiaobao,* to more than 300 million users, providing services such as fund evaluation, market analysis and portfolio management. In 2023, the platform advertised the first generative AI-powered financial assistant in China, *Ma Xiaocai*, to selected users, and the full deployment was conducted in 2024. Compared to the robo-advisor *Zhi Xiaobao*, *Ma Xiaocai* employs Ant's proprietary generative-AI model that understands complex investor queries, synthesizes information from multiple sources, performs context reasoning, and generates personalized responses through interactions. By 2025, *Ma Xiaocai* accumulated more than 70 million active users, making it one of the largest professional AI financial applications in China by user base. Figure 1 illustrates a detailed image on the entrance and content of *Ma Xiaocai*.

2.2 Data Collection

Our sample period is from July 2024 to October 2025. For this sample period, the rollout of *Ma Xiaocai* has been completed, and its adoption has reached a sufficiently large group of

[1] Ant Fortune serves more than 500 million investors, including both money market fund investors, such as *Yu'E Bao* users, and non-money-market mutual fund investors.

investors. Our datasets cover 400,000 randomly selected active mutual fund investors.[2] Active mutual fund investors refer to those who have traded at least once during the sample period or hold more than RMB 100 in non-money-market mutual funds, following D'Acunto, Prabhala, and Rossi, 2019, Bianchi and Briere (2026), Hong, Lu and Pan. (2025) and Ge, Tan, Wu and Zhang. (2022). Altogether Ant platform has more than 70 million active mutual fund investors in our sample period, and our random sample covers approximately 0.6% of the total population.

2.3 Summary Statistics

Summary statistics of the user characteristics are reported in Table 1 Panel A. Notice that these accounts are directly linked to unique national IDs, one user can only have one valid account. Out of the 400,000 active users, 31% of investors have interacted with AI assistant at least once during the sample period. This adoption rate is comparable in magnitude to the rapid diffusion of general-purpose generative AI documented in the U.S., where 39% of adults aged 18-64 reported using generative AI in 2024 (Bick, Blandin, and Deming, 2026). Investors submit an average of 1.68 inquiries to the AI assistant, 52% of which are seeking information. About 56% of investors do not hold a bachelor's degree and are defined as low education investors. In terms of family income, 62.1% of investors report annual household income below RMB 50,000, while 18.8% report income above RMB 100,000. The median per capita disposable income in China is RMB 34,707 in 2024. Using the average household size of 2.62 persons reported in the 2020 Population Census implies an approximate household-level income of RMB 91,000, which is comparable to our sample. The average age of investor is 37.9 years, with a balanced gender distribution (49% male), closely resembling the corresponding figures of 38.8 years and 48.8% in the 2020 National Population Census. About 16% of investors reside in one of Tier-1 cities (Beijing, Shanghai, Guangzhou, and Shenzhen). In terms of investors' financial accounts and investment characteristics, the average account value is approximately RMB 40,000, although the median is substantially lower at RMB 3,454,

[2] Compared with prior studies that typically rely on relatively small sample from individual brokerage platforms to study the usage of robo-advisors or on survey-based evidence of generative AI use (D'Acunto, Prabhala and Rossi, 2019; Blankespoor, Croom and Grant, 2026), our data cover a much broader cross-section of retail investors.

suggesting a right-skewed distribution of financial assets. The average portfolio diversification measure is 0.10, with a median of zero, indicating that many investors hold only one type of fund. The average *RiskPreference* score is 2.38. Investors maintain their accounts for an average of 10.5 years. Approximately 13% of investors use automatic investment services, and on average investors have about three years of investment experience.

Panel B of Table 1 presents summary statistics on AI responses and users' follow-up actions. AI assistant provides a response to each of user's inquiries. We present three properties of AI responses: *AIStrength*, *AIReadability*, and *AIRelevance*, which are LLM-based measures capturing the strength of the AI recommendation, how easily the response can be understood by retail investors, and how closely the response addresses the investor's query, respectively. All three variables range from 0 to 1, with higher scores meaning stronger recommendations, greater readability, and higher relevance. Variable *AISatisfaction* is an LLM-based measure of investors' expressed satisfaction with the AI response based on their subsequent queries, taking 1 for satisfaction, -1 for dissatisfaction, and 0 otherwise. For specific AI recommendations, variable *AIRecTrade* measures the direction and strength of the AI recommendation, ranging from -1 to 1, with -1 indicating a strong sell, 0 indicating a hold recommendation, and 1 indicating a strong buy. Its mean is 0.12, suggesting that AI recommendations are, on average, tilted toward buying. It is intriguing to find that the median is 0.0, suggesting that the typical recommendation is to hold rather than to buy or sell. The dummy variable *DUserTrade* measures the direction of trading within 24 hours after a query, taking 1 for a purchase, -1 for a sale, and 0 for no trade. Most *DUserTrade* takes 0, indicating that most user do not trade after the query. *DUserFollowAI* equals one if the investor follows the corresponding AI recommendation within 24 hours, including buying after a buy recommendation, selling after a sell recommendation, or not trading after a hold recommendation. On average, 51% of AI recommendations are followed.

Panel C reports summary statistics for fund characteristics measured at a monthly frequency from July 2024 to October 2025. The average monthly fund return and volatility are 2.32% and 1.12%, respectively. The average fund manages assets worth approximately 1.87

billion yuan. The average fund age is 6.25 years, and the average annual expense ratio is 0.95%. The average quarterly flow is about 10%, during a period of substantial expansion in the mutual fund industry in China.

3. Empirical Results

3.1 Who Uses the GenAI Assistant?

Financial AI assistant offers information and guidance at a near-zero marginal cost compared to traditional human advisors, assisting investors who may struggle to access and process information (Cao, Jiang, Wang and Yang, 2024). It also serves an educational purpose, improving financial literacy for those who need it most (D'Acunto, Prabhala, and Rossi, 2019). The use of financial AI assistant may vary systematically across investors with different demographic and investment characteristics. We examine which investor characteristics are associated with the usage of the AI investment assistant by estimating the following specification:

$$AIUsage_i = a_0 + a_1 DemoChar_i + a_2 InvestChar_i + u_i, \quad (1)$$

where *AIUsage* captures investors' adoption and usage of the AI investment assistant and includes four measures: *DAIUser*, an indicator equal to one if investor *i* uses the AI investment assistant at least once during the sample period and zero otherwise; *NQuery*, the natural logarithm of one plus the total number of AI queries submitted by investor *i*; *NInfoQuery*, the natural logarithm of one plus the number of information-seeking queries; and *NAdviceQuery*, the natural logarithm of one plus the number of advice-seeking queries.

DemoChar includes a set of investor demographic characteristics: *Age*, measured in years; *Male*, an indicator equal to one for male investors and zero otherwise; *DLowEdu*, an indicator equal to one if the investor has less than a bachelor's degree and zero otherwise; *FamilyIncome*, the investor's annual household income category based on the platform's survey, ranging from

1 to 6; and *DTier1City*, an indicator equal to one if the investor resides in one of China's Tier-1 cities[3], and zero otherwise.

InvestChar includes investor investment characteristics. *AccountValue* is the logarithm of the investor's total financial assets held in the account; *Diversification* measures the degree of portfolio diversification, calculated as one minus the Herfindahl index of the investor's fund holdings. We follow Badarinza, Campbell, and Ramadorai (2016) to construct the measure of fund portfolio Herfindahl index. *RiskPreference* captures the investor risk tolerance based on the platform's risk assessment survey, ranging from 0 to 5; *AccountAge* is the logarithm of the number of months since the investor opened the account; *DAutoinvestUser* is an indicator equal to one if the investor participates in the platform's automatic investment service and zero otherwise; and *InvestmentExp* measures the investor's investment experience.

Table 2 reports empirical results of this test for 181,047 investors with complete demographic information. From column (1), AI adoption declines sharply with age. Each additional year of age is associated with a 0.6 percentage point lower probability of using the AI assistant (t=-68.14). Male investors are 1.4 percentage points less likely to adopt the assistant (t=-6.31). Investors without a bachelor's degree are 2.8 percentage points more likely to adopt (t=12.07), an economically meaningful difference relative to the 31% baseline adoption rate. Adoption also declines with household income: each one-category increase in the six-category income scale is associated with a 1.2 percentage point lower adoption probability (t=-10.99), implying a 6-percentage point difference between the lowest and highest income groups. Investors residing in Tier-1 cities are also 1.6 percentage points less likely to adopt the assistant (t=-5.70). Together, these demographic patterns suggest that financial AI is particularly attractive to younger, less educated, lower-income investors, and those residing in Tier-2 and Tier-3 cities, who may face greater frictions in accessing and processing financial information.

Turning to investment characteristics, adoption increases with account value, portfolio diversification, risk preference, and investment experience, while it decreases with account age.

[3] Tier-1 cities refer to Beijing, Shanghai, Guangzhou, and Shenzhen, four of China's largest cities by GDP, following the classification of the National Bureau of Statistics of China.

In particular, a one-unit increase in *AccountValue* is associated with a 3.8 percentage point higher probability of AI adoption (t=85.83). This positive association is consistent with greater information demand among investors with more financial assets at stake, who have stronger incentives to acquire and process investment information (D'Acunto, Prabhala, and Rossi, 2019). The association is particularly strong for participation in automatic investment services: *DAutoinvestUser* is associated with a 20.6 percentage point higher probability of AI adoption (t=66.08). These findings suggest that financial AI also attracts investors who are more actively engaged with investment products and digital financial services.

Column (2) shows similar patterns for overall usage intensity. Less educated and lower-income investors submit more queries, while query activity is also higher among younger investors and those with greater account value, diversification, risk tolerance, investment experience, and participation in automatic investment services. Columns (3) and (4) further show that most of these patterns hold for both information-seeking and advice-seeking queries. In particular, less educated and lower-income investors both seek more information and more advice from the assistant. One notable exception is gender: male investors submit significantly fewer information-seeking queries (-0.052, t=-18.02) but do not differ significantly from female investors in advice-seeking (0.003, t=0.93), suggesting that the gender difference in AI usage is driven primarily by information demand.

This adoption pattern is notable against the existing evidence on technology-assisted advice and general-purpose generative AI. Adopters of traditional robo-advisors tend to be wealthier and more financially sophisticated (D'Acunto, Prabhala, and Rossi, 2019; Rossi and Utkus, 2024), and users of generative AI are generally more educated, and more financially sophisticated (Blankespoor, Croom, and Grant, 2026; Bick, Blandin, and Deming, 2026). In contrast, we find that adoption of financial generative AI assistants is significantly higher among investors with lower education and lower income. This reversed adoption gradient suggests that conversational financial AI distinguishes itself through an accessible, low-cost, natural-language interface that significantly lowers usability barriers, even for individuals with limited financial education and low income.

3.2 How Does AI Assistant Answer?

To understand AI responses, we first examine whether AI responses follow a structured communication pattern. Specifically, we classify each sentence in response according to its primary function: emotional support, factual information, information analysis, or investment advice. We then normalize the position of each sentence within a response and plot the fitted density curves showing the likelihood that each sentence function appears at different relative positions, with higher density values indicating a greater relative likelihood of occurrence. As shown in Figure 2, factual information and analysis are more concentrated in the middle of responses, whereas investment advice and emotional support are more likely to appear at the beginning and the end. The AI-generated response often begins with emotional acknowledgment or a recommendation, then provides supporting information, and finally reiterates or reinforces its advice at the end. This pattern suggests that AI responses tend to follow a "sympathy, information, advice" structure, which combines investment advice with informational reasoning and emotional support in a structured communication format.

On average, factual information accounts for the largest share of sentences, at 39.42%, followed by information analysis at 29.27% and investment advice at 27.40%. Emotional support accounts for the remaining 3.88%. To be more specific, for factual information, fund-related information is the most prevalent category, appearing in 71% of responses, followed by market information (56%) and investor-specific information (44%), whereas stock-specific information appears in only 6% of responses, consistent with the platform operating as a fund-distribution rather than stock-trading service.[4] For the investment-advice component, position proposals (53%) and asset diagnosis (60%) dominate, while direct fund-manager advice is almost absent (0.06%). Among advisory statements that reference a specific fund, hold recommendations dominate, accounting for 68.58%, with buy recommendations at 30.22% and sell recommendations at only 1.19%. The assistant is thus far more likely to counsel investors to maintain existing positions than to recommend active liquidation, a feature that proves

[4] The distribution varies with query intent: advice-seeking queries are more likely to receive fund information (88%) and investor-specific information (67%), indicating that, unlike general-purpose language models, the assistant integrates users' portfolio holdings and platform-specific product information when investors explicitly seek guidance.

central to the performance results in Section 3.6. Finally, among responses containing emotional support, 47.92% provide emotional comfort, 50.57% offer encouragement, and 1.51% contain both.

The most important component for AI assistant response is its recommendation, which leads to two natural questions: Does AI assistant individualize its responses, according to user's properties and fund properties? Do users actually follow AI assistant recommendations? In this section, we mainly study the first question. Compared to traditional robo-advisors and general-purpose AI, financial generative AI assistants can process richer investor-specific and fund-specific information and generate recommendations through interactive, context-dependent reasoning. We therefore expect generative AI investment assistants to provide customized recommendations based on both investor profiles and the characteristics of the underlying assets.

We examine whether AI assistant's recommendations are customized to user and fund characteristics by estimating the following specification:

$$AIRecTrade_{i,j,k,t} = b_0 + b_1 UserChar_{i,t-1} + b_2 FundChar_{j,t-1} + b_3^{'} Controls_{i,j,k,t} + FE_t + u_{i,j,k,t}. \quad (2)$$

Here the dependent variable, $AIRecTrade_{i,j,k,t}$, captures both the direction and strength of the AI-generated fund recommendation for user *i* on fund *j* in query *k* on day *t*, ranging from -1 to 1. A value of -1 indicates a strong sell recommendation, 0 indicates a hold recommendation, and 1 indicates a strong buy recommendation. Between -1 and 0, a more negative number indicates a stronger sell recommendation, and between 0 and 1, a larger number indicates a stronger buy recommendation. We identify the recommendation's direction and strength using an LLM-based classification of the AI-generated response.

The first set of independent variables, *UserChar*, includes key user characteristics measured at the end of the previous month, including risk preference (*RiskPreference*), which captures the investor's willingness to bear investment risk; the natural logarithm of account value (*AccountValue*), which captures the investor's financial assets; portfolio diversification

(*Diversification*), which captures the degree of investor portfolio diversification; the proportion of equity funds in the portfolio (*EquityFundRatio*), which captures the investor's exposure to equity funds; previous-month portfolio return (*L1mUserRet*), which captures portfolio performance; previous-month portfolio volatility (*L1mUserVolatility*), which captures portfolio risk; and previous-month turnover (*L1mUserTurnover*), which captures recent investor trading intensity. If the AI assistant provides customized recommendations, we expect its recommendations to vary systematically with these investor characteristics.

The second set of independent variables are fund level key variables, *FundChar*, including previous-day return (*L1dFundRet*) and previous-month return (*L1mFundRet*), which capture recent fund performance; previous-month volatility (*L1mFundVolatility*), which captures recent fund risk; the natural logarithm of assets under management (*FundSize*), which captures fund size; the natural logarithm of fund age (*FundAge*), which captures fund maturity; expenses (*Expenses*), which capture the cost of holding the fund; and previous-quarter fund net flow (*L1qFundNetFlow*), which captures recent investor demand for the fund and the flow pressures known to affect fund behavior (Chen, Goldstein, and Jiang, 2010). If the AI assistant incorporates fund-specific information into its recommendations, we expect its recommendations to vary systematically with these fund characteristics.

All specifications control for the query's stated trading intent, including *BuyIntent*, *SellIntent*, and *HoldIntent*, which are indicator variables equal to one if the investor expresses an intention to buy, sell, or hold, respectively, in the query. We also include day fixed effects, so the coefficients capture how AI recommendations vary across investors and funds, holding constant the investor's stated trading intent and common time-specific factors.

The results are reported in Table 3. Column (1) only includes user level information. We find that the AI assistant provides more conservative recommendations to investors with larger accounts (-0.002, t = -6.77), consistent with more cautious position management when more financial assets are at stake. A one-standard-deviation increase in *AccountValue* is associated with an approximately 0.006 unit decrease in *AIRecTrade*. In contrast, more risk-tolerant investors receive more buy-oriented recommendations (0.004, t = 9.15), consistent with their

greater willingness to bear investment risk. The assistant also provides more conservative advice for investors with equity-heavy portfolios (-0.028, t = -12.98) and more volatile recent portfolio returns (-0.016, t = -12.18). This pattern is consistent with suitability-based conditioning: the assistant moderates its advice for investors already carrying substantial risk. The positive coefficient on turnover (t = 6.44) indicates more action-oriented advice to investors who trade frequently: a pattern we revisit when interpreting the following results, since frequent traders are also the investors least likely to follow the advice they receive.

Column (2) includes only fund-level characteristics. Previous-day fund returns are not significantly related to AI recommendations, while higher previous-month returns are associated with more buy-oriented recommendations (0.001, $t = 4.84$), suggesting some return-chasing tendency. The assistant also recommends against volatile funds (-0.025, t = -14.64) and large funds (-0.001, t = -2.11). The strongest single fund determinant is expenses: a one-standard-deviation increase in the expense ratio (0.45 percentage points) shifts the recommendation by roughly 0.024 units toward sell, making fee aversion the assistant's most pronounced advisory tilt. The fee and volatility aversion align the advice with investor interests and contrast with evidence that conflicted human advisors steer clients toward expensive products (Mullainathan, Noeth, and Schoar, 2012; Linnainmaa, Melzer, and Previtero, 2021); the return-chasing tilt mirrors a familiar retail bias and previews the return-predictability results in Section 3.5.

In column (3), we include both user- and fund-level characteristics, and the results remain similar to those in columns (1) and (2). Taken together, our evidence suggests that the financial generative AI assistant systematically customizes its recommendations to both investor profiles and fund characteristics. This contrasts with traditional human advisors, whose advice exhibits limited customization to client characteristics (Foerster, Linnainmaa, Melzer, and Previtero, 2017). While client interaction enables robo-advisors to dynamically incorporate investor risk preferences (Capponi, Ólafsson, and Zariphopoulou, 2022), financial generative AI assistant potentially strengthens personalization by jointly incorporating rich investor- and product-specific information into context-dependent recommendations. Our findings thus highlight

personalization as a distinctive feature of generative AI-based financial advice.

3.3 Do Investors Trade Following the Advice?

Having established that AI recommendations are customized based on user- and fund-level information, we next examine whether investors actually follow these recommendations. Cheng, Lin and Zhao (2025) document a significant decline in stock trading volume during ChatGPT outages, suggesting that access to generative AI can affect investors' trading behavior. Building on this evidence, we hypothesize that AI recommendations influence users' subsequent trading decisions and, more specifically, that investors tend to trade in the direction recommended by AI. We therefore estimate the following regressions to test this hypothesis:

$$DUserTrade_{i,j,k,t\to t+1} = c_0 + c_1 AIRecTrade_{i,j,k,t} + c_2' Controls_{i,j,k,t} + FE_t + u_{i,j,k,t}. \quad (3)$$

The dependent variable, $DUserTrade_{i,j,k,t\to t+1}$, captures investor i's trading direction for fund j after query k submitted on day t, based on the investor's trading activity within 24 hours following the query. It equals 1 if the investor buys the fund, -1 if the investor sells the fund, and 0 if the investor does not trade the fund. The main independent variable, $AIRecTrade_{i,j,k,t}$, is the AI-generated recommendation for fund j in query k submitted on day t. The user-level controls include risk preference, account value, portfolio diversification, equity fund holdings, and past portfolio performance and trading activity; the fund-level controls include past performance and volatility, size, age, expenses, and net flows; and the query-level controls include the user's trading intention. All specifications include day fixed effects. If our hypothesis that investors follow AI advice is correct, coefficient c_1 should be significantly positive.

Table 4 presents the results on whether investors trade on AI advice. Column (1) does not include any controls. The coefficient on *AIRecTrade* is 0.163 (t = 20.20), indicating a strong positive unconditional relation between AI recommendations and investors' subsequent trading decisions. Investors are more likely to buy when the AI recommendation is more positive and more likely to sell when it is more negative.

Column (2) adds the full set of investor-, fund-, and query-level controls. The coefficient on *AIRecTrade* decreases substantially to 0.026 (t = 3.81) but remains positive and statistically significant. The attenuation suggests that part of the unconditional relation in column (1) reflects the assistant's tailoring of recommendations to investor and fund characteristics, as documented in Table 3. Nevertheless, the significant coefficient after controlling for these characteristics indicates that investors' subsequent trades remain directionally aligned with AI recommendations.

Column (3) further includes day fixed effects to absorb common time-varying factors that may jointly affect AI recommendations and investor trading. The coefficient remains stable at 0.024 (t = 4.01), suggesting that the relation is robust to controlling for aggregate time effects. Economically, a one-unit increase in *AIRecTrade* is associated with a 2.4 percentage point increase in the investor's net propensity to buy the queried fund within 24 hours. Overall, the results support our hypothesis that investors follow AI investment advice.

3.4 Who Follows AI Advice?

Notice that 51% of AI recommendations are followed, while 49% are not. What leads to the differences? There are two potential mechanisms: user characteristics and AI response features. For investor characteristics, lower-education investors may have greater difficulty independently evaluating financial information and therefore rely more on external advice. Consistent with this argument, Stolper and Walter (2017) show that education is positively associated with financial literacy and that investors with lower financial literacy are more likely to follow financial advisors' recommendations. Separately, investors who actively seek investment advice may have a stronger demand for guidance and thus be more receptive to the advice they receive. Landis, Fisher and Menges (2022) similarly show that individuals are more likely to use advice when it is solicited rather than unsolicited. We therefore hypothesize that lower-education investors and advice-seeking investors are more likely to follow AI recommendations.

For features of AI-generated responses, beyond the direction of a recommendation itself, the strength of the supporting arguments may shape investors' responses. Asquith, Mikhail and

Au (2005) show that investors respond not only to the direction of analyst recommendations but also to the strength of the arguments supporting them. Following this intuition, we hypothesize that AI recommendations supported by stronger arguments are more likely to be followed by investors. The quality of an AI response may also affect investors' willingness to follow its recommendation. Rennekamp (2012) shows that more readable disclosures are perceived as more credible and generate stronger investor reactions. In the context of generative AI, Ecker, Li, Li and Wu (2026) show that higher AI response quality increases investors' subsequent AI usage, suggesting that users respond to the quality of AI-generated information. We therefore hypothesize that AI recommendations associated with higher-quality responses are more likely to be followed by investors.

We estimate the following regression for these potential mechanisms:

$$DUserFollowAI_{i,j,k,t\to t+1} = d_0 + d_1 UserChar_{i,j,k,t} + d_2 AIChar_{i,j,k,t} + d_3{}' Controls_{i,j,k,t} + FE_t + u_{i,j,k,t}. \quad (4)$$

The dependent variable, $DUserFollowAI_{i,j,k,t}$, captures whether investor i follows the AI recommendation for fund j, which is referenced in query k submitted on day t, based on the investor's trading activity within 24 hours following the query. It equals one if the investor buys after receiving a buy recommendation, sells after receiving a sell recommendation, or does not trade after receiving a hold recommendation, and zero otherwise. $UserChar$ includes *DLowEdu* and *DAdviceSeeking*. *DAdviceSeeking* equals one if the investor's primary query intent is to seek investment advice rather than pursue other objectives, such as seeking information. For AI-response features $AIChar$, *AIStrength* measures the strength of the AI-generated recommendation, ranging from 0 to 1, with higher values indicating that the AI expresses a more explicit and forceful recommendation to take the suggested trading action, and it is computed as the absolute value of the recommendation tone. The other three variables (*AIReadability*, *AIRelevance*, *AISatisfaction*) are proxies for response quality. *AIReadability* is an LLM-generated measure of the linguistic readability of the AI response, with higher values indicating responses that are easier to understand. *AIRelevance* is an LLM-generated measure

of how relevant the AI response is to the investor's query, with higher values indicating better alignment between the response and the question. *AISatisfaction* measures the investor's expressed satisfaction with the AI response, taking the value of 1 for satisfied, -1 for dissatisfied, and 0 if no satisfaction or dissatisfaction is expressed. Control variables, including user-, fund- and query-level characteristics, and fixed effects are specified as in Section 3.3. If our hypothesis that investors with lower education and those who actively seek investment advice are more likely to follow AI advice is correct, we expect d_1 to be positive for *DLowEdu* and *DAdviceSeeking*. Similarly, if investors are more likely to follow AI recommendations that are expressed more strongly and accompanied by higher-quality responses, we expect d_2 to be positive for *AIStrength, AIReadability, AIRelevance, and AISatisfaction*.

Table 5 reports the results on the determinants of investors' decisions to follow AI advice. Panel A examines investor characteristics. Column (1) first examines investors' education. The coefficient on *DLowEdu* is 0.012 (t = 5.21), indicating that investors without a bachelor's degree are 1.2 percentage points more likely to follow AI advice than those with a bachelor's degree or above. Column (2) examines whether investors who actively seek investment advice are more receptive to AI recommendations. The coefficient on *DAdviceSeeking* is 0.043 (t = 12.75), indicating that investors whose primary query intent is advice-seeking are 4.3 percentage points more likely to follow AI advice than investors pursuing other objectives, such as seeking information. This result provides strong support for the hypothesis that investors with lower education and those who actively seek advice are more likely to follow the recommendations they receive. Column (3) includes these investor characteristics jointly. The coefficient on *DLowEdu* remains positive and statistically significant at 0.012 (t = 5.10). Meanwhile, the coefficient on *DAdviceSeeking* remains positive and highly significant and increases slightly to 0.046 (t = 11.91).

Among the controls, two additional patterns stand out. High turnover investors are significantly less likely to follow AI advice (t = -9.35 in Column 3), indicating that the investors whose frequent trading the advice could most usefully restrain are also the least receptive to it, a tension we return to when interpreting the performance results. More diversified investors

are also significantly less likely to follow AI advice (t = -5.30 in Column 3), consistent with more sophisticated investors treating the AI assistant as one source of information among several rather than as an advisor whose recommendations should be followed directly.

Panel B examines features of AI-generated responses. Column (1) first examines recommendation strength. The coefficient on *AIStrength* is 0.041 (t = 6.89), indicating that investors are more likely to follow AI recommendations when the recommendations are expressed more strongly. This result is consistent with the hypothesis that stronger AI recommendations are more likely to be followed.

Columns (2) through (4) examine different dimensions of response quality separately. In Column (2), the coefficient on *AIReadability* is 0.068 (t = 4.37), indicating that investors are more likely to follow AI advice when the response is easier to understand. In Column (3), the coefficient on *AIRelevance* is 0.078 (t = 5.09), showing that recommendations contained in responses that are more closely aligned with investors' queries are more likely to be followed. In Column (4), the coefficient on *AISatisfaction* is 0.012 (t = 2.93), indicating that investors who express greater satisfaction with an AI response are also more likely to follow its recommendation. These results consistently support the hypothesis that AI recommendations associated with higher-quality responses are more likely to be followed.

Column (5) includes recommendation strength and all three response-quality measures jointly. All four coefficients remain positive and statistically significant. The coefficient on *AIStrength* is 0.040 (t = 6.79), while the coefficients on *AIReadability*, *AIRelevance*, and *AISatisfaction* are 0.044 (t = 2.78), 0.060 (t = 3.81), and 0.013 (t = 2.97), respectively. The stability and significance of these coefficients in the joint specification indicate that the effects of recommendation strength and response quality are largely independent of one another. Economically, a one-standard-deviation increase in recommendation strength (0.31) increases the likelihood of following AI advice by approximately 1.2 percentage points. Similarly, one-standard-deviation increases in readability (0.06) and relevance (0.05) each increases the likelihood of following by approximately 0.3 percentage points. Investor's expressed satisfaction with the AI response also significantly predicts following.

Two broader interpretations of the results deserve emphasis. First, following is a choice that responds to the message rather than a mechanical artifact of the follow definition: holding constant investor, fund, and query characteristics as well as day fixed effects, better-crafted and more forceful responses are more likely to be followed. Second, the pattern is consistent with the persuasion evidence for human financial communication. Stronger arguments amplify market reactions to analyst recommendations (Asquith, Mikhail, and Au, 2005), while accessible messages travel further under limited attention (Peng and Xiong, 2006). Our results extend this intuition to human–AI interactions by showing that individual investors' decisions to follow AI advice systematically vary with observable characteristics of the AI-generated response.

3.5 Do AI Recommendations Predict Fund Returns?

We next examine whether AI recommendations can predict future fund returns. If AI recommendations incorporate value-relevant information that is not yet fully reflected in fund prices, more favorable recommendations should predict higher subsequent returns. Conversely, if the AI primarily synthesizes publicly available information that is already incorporated into prices, its recommendations may have little predictive power for future returns. To examine the predictive content of AI recommendations, we estimate the following regressions:

$$FundPerformance_{i,j,k,t+1\to t+h} = e_0 + e_1 AIRecTrade_{i,j,k,t} + e_2' Controls_{i,j,k,t} + FE_t + u_{i,j,k,t}. \quad (5)$$

The dependent variable $FundPerformance_{i,j,k,t+1\to t+h}$ includes the next-day, next-week and next-month fund j returns following query k submitted by user i on day t. The main independent variable, $AIRecTrade_{i,j,k,t}$, captures both the direction and strength of the AI-generated recommendation for fund j. Controls include past fund returns over multiple horizons, past fund volatility, fund size, fund age, expenses, and past fund net flows, as well as the query's stated trading intent. All specifications include day fixed effects, and standard errors are double-clustered at the fund and day levels.

As shown in Table 6, we find no evidence that AI recommendations predict future fund

returns at any horizon. The coefficients on *AIRecTrade* are 0.036 (t = 0.60) over the next day, 0.109 (t = 0.68) over the next week, and -0.455 (t = -1.37) over the next month. The estimates are economically small as well as statistically insignificant: taking the next-month point estimate at face value, a one-standard-deviation more favorable recommendation (0.20) is associated with a 9-basis point lower next-month return, and the short-horizon estimates are an order of magnitude smaller still.

Controlling for the fund's return history is particularly important for interpreting these results. As shown in Table 3, the AI assistant tends to issue more favorable recommendations for funds with stronger past returns. Without adequately controlling for past performance, a naïve specification could therefore confound the predictive content of AI recommendations with return dynamics following recent fund performance. By conditioning on the fund's prior returns over the past day, month, and three months, our specification isolates the incremental information contained in the AI recommendation beyond information already reflected in the fund's recent return history. The results indicate that AI recommendations contain no significant incremental information about subsequent fund returns.

Appendix Table IA4 provides further evidence using a fund-day design. We aggregate AI recommendations to the fund-day level and estimate Fama-MacBeth regressions of future fund returns on the fund's buy ratio. The estimated coefficients are uniformly insignificant, leading to the same conclusion. Taken together, the evidence suggests that whatever value AI advice may provide to investors, it does not appear to arise from identifying funds that subsequently outperform.

3.6 Following AI Advice and Portfolio Outcomes

3.6.1 Investment Performance

We first examine whether following AI recommendations affects investors' subsequent investment performance. On the one hand, AI advice may help investors make better investment decisions by providing systematic and disciplined guidance. D'Acunto, Prabhala and Rossi (2019) show that robo-advising can mitigate investors' behavioral biases and

improve investment performance, suggesting that technology-based financial advice can help investors overcome behavioral frictions that otherwise lead to suboptimal decisions. Following this intuition, we hypothesize that following AI recommendations improves investors' subsequent investment performance. On the other hand, following AI advice need not improve investment outcomes. LLM-generated recommendations may themselves reflect behavioral biases or other systematic distortions. Bini, Cong, Huang and Jin (2026) show that LLMs exhibit behavioral biases, raising the possibility that investors who follow their recommendations may inherit or reinforce these biases rather than correct them. We therefore propose the competing hypothesis that following AI recommendations worsens investors' subsequent investment performance. To investigate investors' portfolio outcomes after following AI recommendations, we construct the following regression:

$$Performance_{i,j,k,t+1\to t+h} = f_0 + f_1 * DUserFollowAI_{i,j,k,t} + f_2' Controls_{i,j,k,t} + FE_t + u_{i,j,k,t}, \quad (6)$$

where h corresponds to one day, one week, or one month. The dependent variable, $Performance_{i,j,k,t+1\to t+h}$, measures investor i's portfolio performance over the corresponding horizon following query k submitted on day t. The main independent variable, $DUserFollowAI_{i,j,k,t}$, captures whether investor i follows the AI recommendation for fund j, which is referenced in query k submitted on day t, based on the investor's trading activity within 24 hours following the query. Controls including user-, fund-, and query-level characteristics and day fixed effects are as specified in Section 3.3. If following AI advice improves investors' subsequent performance, we expect f_1 to be positive for portfolio returns.

Table 7 shows the association between following AI recommendations and investors' subsequent portfolio returns. Investors who follow AI advice earn significantly higher subsequent portfolio returns at every horizon. Following an AI recommendation is associated with 1.8 basis points higher returns over the next day (t = 2.83), 7.3 basis points over the next week (t = 3.98), and 10.9 basis points over the next month (t = 3.33), with the one-month estimate corresponding to approximately 1.3% on an annualized basis. These results are

consistent with the hypothesis that following AI recommendations is associated with better subsequent investment performance.

To understand which type of recommendation drives this improvement in portfolio returns, we further decompose *DUserFollowAI* into *DBuyFollowAI*, *DSellFollowAI*, and *DHoldFollowAI*, indicating whether the investor follows a buy, sell, or hold recommendation, respectively. The decomposition reveals that the positive association is concentrated in following hold recommendations. Investors who receive a hold recommendation and refrain from trading earn 2.4 basis points higher returns over the next day, 9.5 basis points over the next week, and 17.5 basis points over the next month, with the one-month effect highly significant (t = 4.40). In contrast, following a buy recommendation is not associated with higher returns over the next day or week and is associated with 15.9 basis points lower returns over the next month (t = -2.37). Following sell recommendations has no statistically significant association with subsequent returns at any horizon.

This decomposition helps clarify what drives the overall finding that following AI advice improves investor performance. As shown in Table 6, more favorably recommended funds do not subsequently outperform, suggesting that the higher portfolio returns associated with following AI advice are unlikely to arise from superior fund-selection ability. Consistent with this evidence, the positive association in Table 7 is driven by investors who follow hold recommendations rather than those who act on buy or sell recommendations. The value associated with following AI advice therefore appears to be concentrated precisely where the AI counsels inaction. Rather than helping investors identify funds that subsequently outperform, AI advice may create value by discouraging potentially detrimental trades and inducing greater trading restraint.

3.6.2 Volatility, Sharpe Ratios, and Turnover

We next examine whether the value-enhancing evidence extends beyond portfolio returns to other portfolio outcomes. Specifically, Table 8 examines investors' subsequent portfolio volatility, Sharpe ratios, and turnover following each query. Portfolio volatility and turnover

are expressed in percentage terms for ease of interpretation. As in Table 7, we first examine overall following behavior and then decompose *DUserFollowAI* by the type of recommendation followed.

Panel A of Table 8 shows that following AI advice is not significantly associated with portfolio volatility over either the subsequent week or month (t = 1.12 and 1.26, respectively). In contrast, investors who follow AI advice achieve significantly higher risk-adjusted performance. Following is associated with a 0.006 increase in the Sharpe ratio over the subsequent week (t = 2.46) and a 0.002 increase over the subsequent month (t = 1.97). The largest association appears in portfolio turnover. Investors who follow AI advice have turnover that is 10.5 percentage points lower over the subsequent week (t = -28.09) and 13.6 percentage points lower over the subsequent month (t = -22.42). Relative to the sample average monthly turnover of 65.49%, the latter represents a reduction of approximately 21%. Together with the higher portfolio returns documented in Table 7, the higher Sharpe ratios and substantially lower turnover provide further evidence that following AI advice improves investor performance.

To understand which type of recommendation accounts for these patterns, Panel B of Table 8 decomposes *DUserFollowAI* into *DBuyFollowAI*, *DSellFollowAI*, and *DHoldFollowAI*, indicating whether the investor follows a buy, sell, or hold recommendation, respectively. Following buy or sell recommendations is associated with higher subsequent turnover, consistent with the fact that following these recommendations requires investors to trade. In contrast, the reduction in overall turnover is concentrated among investors who follow hold recommendations. Their portfolio turnover decreases by 16.4 percentage points over the subsequent week and 20.1 percentage points over the subsequent month. Hold-followers are also the only group that exhibits significantly higher Sharpe ratios, with increases of 0.009 over the subsequent week (t = 2.85) and 0.004 over the subsequent month (t = 2.88). Although following hold recommendations is associated with a small increase in measured portfolio volatility (0.015, t = 2.74), the improvement in risk-adjusted performance indicates that the accompanying return improvement more than offsets the increase in volatility. The small volatility increase is itself unsurprising: refraining from trading means forgoing the mechanical

de-risking that portfolio rebalancing into cash would otherwise provide.

Taken together, Tables 6 through 8 provide consistent evidence that following AI advice improves investor performance. Table 6 shows that AI recommendations do not predict subsequent fund returns, suggesting that the AI does not create value by identifying funds that subsequently outperform. Tables 7 and 8, however, show that investors who follow AI advice subsequently earn higher portfolio returns and Sharpe ratios while trading substantially less. The decomposition further shows that these improvements are concentrated among investors who follow hold recommendations. Thus, the value associated with following AI advice appears to operate primarily through trading discipline rather than superior fund selection: by encouraging investors to refrain from trading, AI advice may help curb excessive or poorly timed trades, consistent with the evidence on the costs of excessive trading among retail investors in Barber and Odean (2000).

The economic magnitude of the turnover reduction is also meaningful. Using a round-trip trading cost of 89 basis points, the 13.6-percentage-point reduction in monthly turnover implies transaction-cost savings of approximately 12 basis points per month. This magnitude is comparable to the 10.9 basis point improvement in monthly portfolio returns documented in Table 7, suggesting that reduced trading costs provide a plausible economic channel through which following AI advice is associated with better portfolio performance.

3.7 Response Quality and Further Engagement

Finally, we examine the factors associated with investors' further engagement with the AI assistant by estimating the following regression:

$$FutureUse_{i,k,t\to t+30} = g_0 + g_1 * Char_{i,k,t} + g_2' Controls_{i,k,t} + FE_t + u_{i,k,t}, \quad (7)$$

where the dependent variable, $FutureUse_{i,k,t\to t+30}$, measures investor i's subsequent engagement with the AI assistant following query k submitted on day t, defined as the natural logarithm of one plus the number of additional queries submitted by the investor within the following month. The key explanatory variables capture both investor characteristics and

characteristics of the AI-generated response. Investor characteristics include *DLowEdu*, *DAdviceSeeking*. Response characteristics include *AIStrength*, *AIReadability, AIRelevance*, and *AISatisfaction*. The user-level controls include risk preference, account value, portfolio diversification, equity fund holdings, and past portfolio performance and trading activity and the query-level controls include the user's trading intention.

Table 9 reports the results. Columns (1) through (6) introduce the explanatory variables separately, and the final column includes them jointly. We first examine investor characteristics. Investors without a bachelor's degree exhibit significantly greater subsequent AI usage: the coefficient on *DLowEdu* is 0.335 (t = 5.44), corresponding to roughly 40% more queries over the following month. Advice-seeking investors exhibit greater subsequent engagement as well. The coefficient on *DAdviceSeeking* is 0.070 (t = 3.00), indicating that investors who actively seek investment advice subsequently use the AI assistant more than those pursuing other objectives, such as seeking information. Together, these results suggest that investors with lower education, as well as those with a stronger demand for investment guidance, tend to engage more intensively with the AI assistant over time.

We next examine features of AI-generated responses. Recommendation strength is positively but only marginally associated with subsequent usage (0.038, t = 1.80), providing some evidence that more strongly expressed recommendations encourage further engagement. Readability, in contrast, does not significantly predict subsequent usage (-0.144, t = -1.31), even though it significantly predicts whether investors follow AI advice in Table 5. Relevance is strongly associated with subsequent engagement (0.152, t = 4.14): investors return more when the AI response is more closely aligned with their query. Finally, investors' expressed satisfaction with the response also strongly predicts subsequent usage (0.432, t = 5.60).

The final column includes all investor and response characteristics jointly, and the main results remain broadly consistent. Investors without a bachelor's degree submit significantly more subsequent queries (0.329, t = 5.49), and advice-seeking investors also return more (0.043, t = 1.72). Among response characteristics, recommendation strength becomes insignificant (0.021, t = 0.88), as does readability (−0.163, t = −1.29), while relevance remains significantly

associated with subsequent usage (0.101, t = 2.41), as does investors' expressed satisfaction (0.429, t = 4.66). The satisfaction effect is economically large: expressing satisfaction with a response is associated with approximately 54% more queries over the following month relative to a neutral reaction.

These results carry two broader messages. First, further engagement appears to respond more strongly to the substance of AI responses than to their presentation. While readability significantly predicts whether investors follow an AI recommendation, it does not predict whether they subsequently return to the assistant; relevance, by contrast, matters for both. Appendix Table IA5 provides a complementary pattern for immediate satisfaction, which responds to readability. Taken together, presentation may shape investors' immediate responses to AI-generated content, whereas substantive relevance is more important for sustained engagement. Second, further engagement exhibits the same financial-literacy gradient documented for adoption and advice following: investors with lower education are more likely to adopt the AI assistant, more likely to follow its recommendations, and more likely to continue using it after their initial interactions. AI-based financial assistance therefore appears particularly salient for less financially sophisticated investors, who not only rely more heavily on AI recommendations but also sustain their engagement with the assistant over time.

4. Conclusion

Generative AI assistants are rapidly becoming a standard feature of retail investment platforms, yet until now there has been little direct evidence on how households actually use them. Using account-level data on 400,000 randomly selected mutual fund investors on a large Chinese fintech platform, with detailed records linking investor queries, AI-generated responses, and subsequent trading, we provide the first large-scale field evidence on AI financial advice in organic deployment. Our data allow us to trace the complete chain of the advisory interaction: who adopts the assistant, what investors ask, what the AI recommends and why, whether investors follow the advice, and whether following it improves investment outcomes.

Our findings paint a consistent picture across every stage of this chain. Adoption

concentrates among investors with lower education and lower income, precisely the groups least served by traditional advisory channels, reversing the sophistication gradient documented for robo-advisors and general-purpose AI tools. The assistant's recommendations are genuinely customized: they condition on investor risk profiles and portfolio composition, screen strongly against high-fee and high-volatility funds, and are dominated by counsel to hold rather than to trade. The advice moves behavior, with trades in the following 24 hours tilting significantly toward the recommended direction, and it is more likely to be followed when the recommendation is stronger and the response is more relevant and readable. Most strikingly, the advice creates investment value without any fund-selection skill: recommended funds do not subsequently outperform, yet investors who follow the advice earn higher portfolio returns and Sharpe ratios at essentially unchanged risk.

The resolution of this apparent paradox is the paper's central insight. The gains from following AI advice are concentrated mainly among investors who follow hold recommendations, and following the advice reduces portfolio turnover by roughly one-fifth relative to the sample mean, generating transaction-cost savings of the same order as the observed return improvement. The economic value of AI financial advice in our setting therefore operates through trading discipline rather than information: the assistant helps investors not by telling them what to buy, but by persuading them, at scale and at near-zero marginal cost, not to trade. This mechanism connects two literatures that have developed separately, the long-standing evidence that retail investors trade too much at great cost and the emerging evidence on the capabilities of large language models, and suggests that the most valuable output of an AI advisor may be restraint rather than prediction.

These findings carry implications for platforms, regulators, and the broader debate over the economic value of generative AI. For platforms, the results suggest that the welfare-relevant dimensions of an AI assistant are not forecasting accuracy but response relevance, communication quality, and the willingness to counsel inaction, features that also sustain user engagement. Notably, the platform's incentives are not obviously aligned with this channel, since an assistant that suppresses trading may reduce fee revenue; that the deployed system

nonetheless counsels holding in more than two-thirds of fund-specific advice is itself an institutional fact worth understanding. For regulators, the concentration of adoption and reliance among less financially literate investors cuts both ways: AI advice reaches households that human advisors do not, but the same households are least equipped to evaluate the advice they receive, so the quality, incentives, and accountability of deployed advisory models warrant scrutiny commensurate with their reach. For the broader debate, our evidence indicates that generative AI can create measurable value in household finance, and that the value need not come from machine intelligence outperforming markets: in our setting it comes from correcting well-documented human error.

References

Asquith, P., Mikhail, M. B., & Au, A. S. (2005). Information content of equity analyst reports. Journal of Financial Economics, 75(2), 245-282.

Badarinza, C., Campbell, J. Y., & Ramadorai, T. (2016). International comparative household finance. Annual Review of Economics, 8, 111-144.

Barber, B. M., & Odean, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. The Journal of Finance, 55(2), 773-806.

Bianchi, M., & Brière, M. (2026). Human-robot interactions in investment decisions. Management Science, 72(1), 14-31.

Bick, A., Blandin, A., & Deming, D. J. (2026). The rapid adoption of generative AI. Management Science.

Bini, P., Cong, L. W., Huang, X., & Jin, L. J. (2026). Behavioral economics of AI: LLM biases and corrections (NBER Working Paper No. 34745). National Bureau of Economic Research.

Blankespoor, E., Croom, J., & Grant, S. M. (2026). Generative AI and investor processing of financial information. Journal of Accounting and Economics, forthcoming.

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

Capponi, A., Olafsson, S., & Zariphopoulou, T. (2022). Personalized robo-advising: Enhancing investment through client interaction. Management Science, 68(4), 2485-2512.

Chatterji, A., Cunningham, T., Deming, D. J., Hitzig, Z., Ong, C., Shan, C. Y., & Wadman, K. (2025). How people use ChatGPT (NBER Working Paper No. 34255). National Bureau of Economic Research.

Chen, J., Tang, G., Zhou, G., & Zhu, W. (2025). ChatGPT and DeepSeek: Can they predict the stock market and macroeconomy? arXiv preprint arXiv:2502.10008.

Chen, Q., Goldstein, I., & Jiang, W. (2010). Payoff complementarities and financial fragility: Evidence from mutual fund outflows. Journal of Financial Economics, 97(2), 239-262.

Chen, Y., Kelly, B. T., & Xiu, D. (2022). Expected returns and large language models. SSRN Working Paper No. 4416687.

Cheng, Q., Lin, P., & Zhao, Y. (2025). Does generative AI facilitate investor trading? Early evidence from ChatGPT outages. Journal of Accounting and Economics, 101821.

Chinco, A., Clark-Joseph, A. D., & Ye, M. (2019). Sparse signals in the cross-section of returns. Journal of Finance, 74(1), 449-492.

Dammon, R. M., Spatt, C. S., & Zhang, H. H. (2004). Optimal asset location and allocation with taxable and tax-deferred investing. Journal of Finance, 59(3), 999-1037.

Dou, W. W., Goldstein, I., & Ji, Y. (2025). AI-powered trading, algorithmic collusion, and price efficiency (NBER Working Paper No. 34054). National Bureau of Economic Research.

Dou, W. W., Goldstein, I., & Ji, Y. (2026). Financial market fragility in the era of AI planning. Working paper, University of Pennsylvania.

Dou, W. W., Goldstein, I., Li, Z., & Yang, L. (2026). Human edge, machine limits: AI-human competition in financial markets. Working paper, University of Pennsylvania.

D'Acunto, F., Prabhala, N., & Rossi, A. G. (2019). The promises and pitfalls of robo-advising. The Review of Financial Studies, 32(5), 1983-2020.

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

Foerster, S., Linnainmaa, J. T., Melzer, B. T., & Previtero, A. (2017). Retail financial advice: Does one size fit all? Journal of Finance, 72(4), 1441-1482.

Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. Journal of Finance, 77(1), 5-47.

Ge, H., Tan, L., Wu, H., & Zhang, X. (2022). How can robot investment assistants help: Collecting information or providing advice? Evidence from China. PBCSF-NIFR Research Paper.

Goldstein, I., Jiang, W., & Karolyi, G. A. (2019). To FinTech and beyond. Review of Financial Studies, 32(5), 1647-1661.

Goldstein, I., Spatt, C. S., & Ye, M. (2021). Big data in finance. The Review of Financial Studies, 34(7), 3213-3225.

Gomes, F., Haliassos, M., & Ramadorai, T. (2021). Household finance. Journal of Economic Literature, 59(3), 919-1000.

He, L., Tan, L., & Zhang, X. (2026). Large language models and return prediction in global capital markets. Working paper, Tsinghua University.

Hong, C. Y., Lu, X., & Pan, J. (2025). FinTech platforms and mutual fund distribution. Management Science, 71(1), 488-517.

Landis, B., Fisher, C. M., & Menges, J. I. (2022). How employees react to unsolicited and solicited advice in the workplace: Implications for using advice, learning, and performance. Journal of Applied Psychology, 107(3), 408-424.

Linnainmaa, J. T., Melzer, B. T., & Previtero, A. (2021). The misguided beliefs of financial advisors. Journal of Finance, 76(2), 587-621.

Lopez-Lira, A., & Tang, Y. (2026). Can ChatGPT forecast stock price movements? Return predictability and large language models. Journal of Financial Economics, forthcoming.

Lu, X., Li, Y., & Shi, F. (2025). Robo-advising meets large language models. Working paper, University of Hong Kong.

Lusardi, A., & Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. Journal of Economic Literature, 52(1), 5-44.

Mullainathan, S., Noeth, M., & Schoar, A. (2012). The market for financial advice: An audit study (NBER Working Paper No. 17929). National Bureau of Economic Research.

Peng, L., & Xiong, W. (2006). Investor attention, overconfidence and category learning. Journal of Financial Economics, 80(3), 563-602.

Rennekamp, K. (2012). Processing fluency and investors' reactions to disclosure readability. Journal of Accounting Research, 50(5), 1319-1354.

Rossi, A. G., & Utkus, S. (2024). The diversification and welfare effects of robo-advising. Journal of Financial Economics, 157, 103869.

Sheng, J., Sun, Z., Yang, B., & Zhang, A. L. (2026). Generative AI and asset management.

Review of Financial Studies, forthcoming.
Stolper, O. A., & Walter, A. (2017). Financial literacy, financial advice, and financial behavior. Journal of Business Economics, 87(5), 581-643.
van Rooij, M., Lusardi, A., & Alessie, R. (2011). Financial literacy and stock market participation. Journal of Financial Economics, 101(2), 449-472.

Figure 1. AI Investment Assistant

This figure presents a snapshot of *Ma Xiaocai*, the AI investment assistant available on the Ant platform. The interface illustrates its key features, including suggested investment-related questions, market insights, portfolio analysis, investment opportunity identification, and an interactive input box through which investors can submit questions by text or voice. English translations are added below the original Chinese text for illustration. Examples of AI queries and corresponding responses are provided in Internet Appendix Table IA3.



Figure 2. Positional Distribution of Sentence Functions within AI Responses

This figure reports the positional distribution of different sentence functions within AI responses. For each response, sentences are classified by their primary function and mapped to their relative position in the response. The figure plots fitted density curves based on the probability of each sentence function occurring at different relative positions, with higher density values indicating a greater relative likelihood of occurrence. Factual informational and information analysis sentences are more likely to appear in the middle of the response, whereas advisory and emotionally supportive sentences are more concentrated at the beginning and the end.

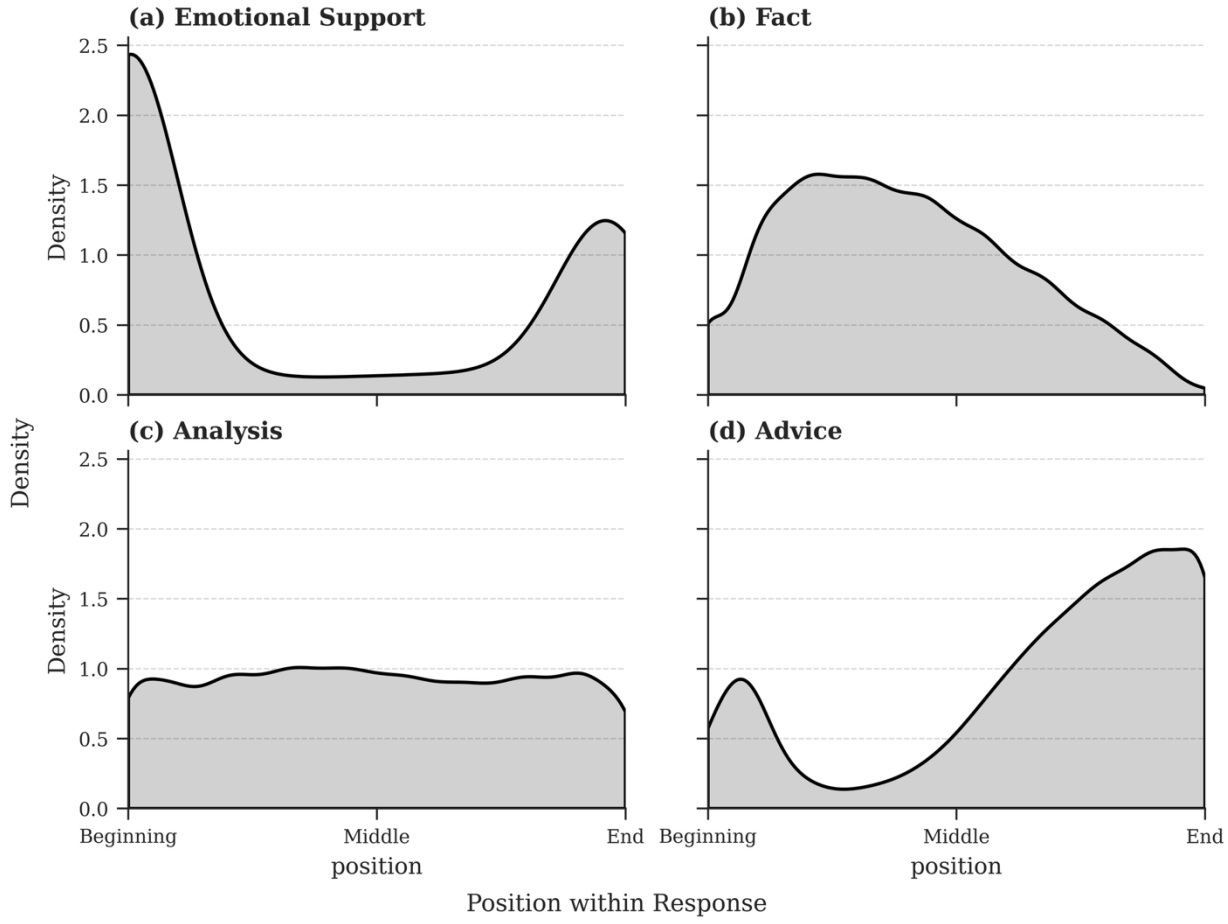


Table 1. Summary Statistics

This table reports the summary statistics for the main variables used in our study. The sample period spans from July 2024 to October 2025. Our sample consists of 400,000 randomly selected active individual investors on the Ant platform who executed at least one fund transaction or held at least RMB 100 in funds during the sample period. Panel A reports investor characteristics, including *DAIUser* (AI assistant usage), *AllInquiries* (total AI queries), *InformationInquiries* (information-seeking queries), *Age* (investor age), *Male* (investor gender), *DLowEdu* (lower educational attainment), *FamilyIncome* (household income), *DTier1City* (Tier-1 city residence), *AccountValue* (total financial assets), *Diversification* (portfolio diversification), *RiskPreference* (risk tolerance), *AccountAge* (years since account opening), *DAutoinvestUser* (auto-invest usage), and *InvestmentExp* (investment experience). Panel B reports AI interactions, including *AIStrength* (AI recommendation strength, ranging from 0 to 1), *AIReadability* (response readability, ranging from 0 to 1), *AIRelevance* (response relevance, ranging from 0 to 1), *AISatisfaction* (investor satisfaction with the response: 1 for satisfied, −1 for dissatisfied, and 0 otherwise), *AIRecTrade* (AI recommendation signal, ranging from -1 to 1), *DUserTrade* (investor's subsequent trading decision: 1 for buy, −1 for sell, and 0 for no trade), and *DUserFollowAI* (whether the investor follows the AI recommendation: 1 for follow and 0 otherwise). Panel C reports fund characteristics, including *FundReturn* (fund return), *FundVolatility* (fund return volatility), *FundSize* (fund total net assets), *FundAge* (fund age), *Expenses* (fund expense ratio), and *FundNetFlow* (fund net flows). For each variable, we report the number of observations, mean, standard deviation, 25th percentile, median, and 75th percentile. Detailed definitions and constructions of all variables are provided in Appendix Table A1.

Panel A. User characteristics

Variables	*Obs.*	*Mean*	*Std*	*P25*	*P50*	*P75*
DAIUser	400000	0.31	0.46	0.00	0.00	1.00
AllInquiries	400000	1.68	9.08	0.00	0.00	1.00
InformationInquiries	400000	0.88	5.51	0.00	0.00	0.00
Age	399999	37.91	12.18	28.00	36.00	46.00
Male	400000	0.49	0.50	0.00	0.00	1.00
DLowEdu	237917	0.56	0.50	0.00	1.00	1.00
FamilyIncome	390332	1.67	1.05	1.00	1.00	2.00
DTier1City	400000	0.16	0.37	0.00	0.00	0.00
AccountValue(CNY)	400000	40400.90	167548.00	383.08	3454.26	24554.28
Diversification	344350	0.10	0.17	0.00	0.00	0.15
RiskPreference	344409	2.38	1.22	2.00	2.00	3.00
AccountAge(Year)	400000	10.49	3.71	7.92	10.17	12.92
DAutoinvestUser	358634	0.13	0.34	0.00	0.00	0.00
InvestmentExp	386454	2.82	1.24	2.00	3.00	4.00

Panel B. AI interaction

Variables	*Obs.*	*Mean*	*Std*	*P25*	*P50*	*P75*
AIStrength	253,740	0.59	0.31	0.38	0.63	0.83
AIReadability	253,740	0.75	0.06	0.70	0.80	0.80
AIRelevance	253,740	0.87	0.05	0.80	0.90	0.90
AISatisfaction	253,740	0.01	0.18	0.00	0.00	0.00
AIRecTrade	253,740	0.12	0.20	0.00	0.00	0.25
DUserTrade	253,740	0.00	0.37	0.00	0.00	0.00
DUserFollowAI	253,740	0.51	0.50	0.00	1.00	1.00

Panel C. Fund characteristics

Variables	*Obs.*	*Mean*	*Std*	*P25*	*P50*	*P75*
FundReturn (%)	101375	2.32	6.56	-0.66	0.49	3.94
FundVolatility (%)	101375	1.12	0.99	0.28	0.93	1.62
FundSize (RMB billion)	101375	1.87	6.95	0.14	0.49	1.66
FundAge (Year)	101375	6.25	4.41	3.05	5.02	8.77
Expenses (%)	101375	0.95	0.45	0.60	1.00	1.40
FundNetFlow	101375	0.10	0.68	-0.11	-0.03	0.06

Table 2. Individual Investor's AI Adoption

This table reports the determinants of individual investors' AI adoption. The sample consists of 400,000 randomly selected active individual investors between July, 2024 and October, 2025. Active investors are defined as users who have executed at least one fund transaction or hold at least 100 RMB in funds on the platform during the sample period. We estimate OLS regressions as specified in equation (1). The dependent variables are *DAIUser*, an indicator equal to one if investor *i* uses the AI investment assistant at least once during the sample period and zero otherwise; *NQuery*, the natural logarithm of one plus the total number of AI queries submitted by investor *i*; *NInfoQuery*, the natural logarithm of one plus the number of information-seeking queries; and *NAdviceQuery*, the natural logarithm of one plus the number of advice-seeking queries. The independent variables are investor characteristics including *age*, gender, education level (*DLowEdu*), residence city (*DTier1City*), annual household income (*FamilyIncome*), account value (*AccountValue*), portfolio diversification (*Diversification*), risk preference (*RiskPreference*), account age (*AccountAge*), *DAutoinvestUser* and investment experience (*InvestmentExp*). Reported values in brackets are t-statistics. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Dep.var	*DAIUser*	*NQuery*	*NInfoQuery*	*NAdviceQuery*
Variables	(1)	(2)	(3)	(4)
Age	-0.006***	-0.008***	-0.005***	-0.005***
	[-68.14]	[-51.59]	[-39.93]	[-38.72]
Male	-0.014***	-0.036***	-0.052***	0.003
	[-6.31]	[-10.02]	[-18.02]	[0.93]
DLowEdu	0.028***	0.074***	0.066***	0.043***
	[12.07]	[18.78]	[20.99]	[14.06]
FamilyIncome	-0.012***	-0.015***	-0.007***	-0.011***
	[-10.99]	[-8.28]	[-5.18]	[-8.06]
DTier1City	-0.016***	-0.047***	-0.036***	-0.034***
	[-5.70]	[-9.88]	[-9.36]	[-9.20]
AccountValue	0.038***	0.066***	0.045***	0.044***
	[85.83]	[89.47]	[76.31]	[76.48]
Diversification	0.255***	0.620***	0.404***	0.499***
	[38.72]	[55.80]	[45.36]	[57.75]
RiskPreference	0.025***	0.048***	0.035***	0.032***
	[23.19]	[25.72]	[23.29]	[21.90]
AccountAge	-0.148***	-0.286***	-0.195***	-0.199***
	[-45.59]	[-52.18]	[-44.55]	[-46.82]
DAutoinvestUser	0.206***	0.481***	0.342***	0.350***
	[66.08]	[91.61]	[81.41]	[85.82]
InvestmentExp	0.038***	0.051***	0.034***	0.030***
	[38.69]	[30.85]	[25.19]	[22.91]
Constant	0.793***	1.256***	0.823***	0.824***
	[56.84]	[53.37]	[43.67]	[45.01]
Adj. R-squared	0.150	0.180	0.138	0.151
Obs.	181,047	181,047	181,047	181,047

Table 3. Properties of AI Responses

This table examines whether AI provides customized fund recommendations based on investor and fund characteristics. We estimate OLS regressions as specified in equation (2). The dependent variable is *AIRecTrade*, which captures both the direction and strength of the AI-generated fund recommendation, ranging from -1 to 1. A value of -1 indicates a strong sell recommendation, 0 indicates a hold recommendation, and 1 indicates a strong buy recommendation. The sign captures the direction of the recommendation, and the absolute value captures its strength. Column (1) examines the relation between investor characteristics and AI recommendations. Investor characteristics are measured at the end of the previous month and include the natural logarithm of account value (*AccountValue*), risk preference (*RiskPreference*), portfolio diversification (*Diversification*), the proportion of equity funds in the investor's portfolio (*EquityFundRatio*), previous-month portfolio return (*L1mUserRet*), volatility (*L1mUserVolatility*), and turnover (*L1mUserTurnover*). Column (2) examines the relation between fund characteristics and AI recommendations. Fund characteristics include previous-day return (*L1dFundRet*), previous-month return (*L1mFundRet*), and volatility (*L1mFundVolatility*), log size of the assets under management (*FundSize*), the log of fund age (*FundAge*), expenses (*Expenses*), and previous-quarter fund net flow (*L1qFundNetFlow*). Column (3) jointly includes both investor and fund characteristics. All regressions control for question-level trading intent, including buy, sell, and hold intent. Day fixed effect is included in all regressions. Standard errors are clustered at day level. Reported values in brackets are t-statistics. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Dep.var	*AIRecTrade*	*AIRecTrade*	*AIRecTrade*
Variables	(1)	(2)	(3)
AccountValue	-0.002***		-0.002***
	[-6.77]		[-6.00]
RiskPreference	0.004***		0.004***
	[9.15]		[9.62]
Diversification	-0.003		0.015***
	[-1.16]		[6.05]
EquityFundRatio	-0.028***		-0.017***
	[-12.98]		[-8.49]
L1mUserRet(%)	0.000		0.000
	[0.62]		[0.60]
L1mUserVolatility(%)	-0.016***		-0.003***
	[-12.18]		[-2.61]
L1mUserTurnover(%)	0.000***		0.000***
	[6.44]		[5.83]
L1dFundRet(%)		0.001	0.001
		[1.36]	[1.46]
L1mFundRet(%)		0.001***	0.001***
		[4.84]	[4.62]
L1mFundVolatility(%)		-0.025***	-0.024***
		[-14.64]	[-13.75]

FundSize		-0.001**	-0.001*
		[-2.11]	[-1.78]
FundAge		-0.002	-0.001
		[-1.42]	[-1.00]
Expenses(%)		-0.061***	-0.058***
		[-27.62]	[-27.36]
L1qFundNetFlow		-0.001***	-0.001***
		[-3.82]	[-4.38]
Constant	0.150***	0.219***	0.219***
	[47.41]	[47.85]	[43.90]
Controls	Yes	Yes	Yes
Day FE	Yes	Yes	Yes
Adj. R-squared	0.072	0.103	0.105
Obs.	253,724	253,724	253,724

Table 4. How do Investors Trade after AI Responses

This table reports the regression results on how AI investment recommendations affect investors' subsequent trading decisions. We estimate the panel regressions as specified in Equation (3). The dependent variable, *DUserTrade*, equals 1 if the investor buys the corresponding fund within 24 hours after the query, −1 if the investor sells the corresponding fund, and 0 otherwise. The main independent variable, *AIRecTrade*, captures both the direction and strength of the AI-generated fund recommendation, ranging from -1 to 1. A value of -1 indicates a strong sell recommendation, 0 indicates a hold recommendation, and 1 indicates a strong buy recommendation. The sign captures the direction of the recommendation, and the absolute value captures its strength. We control for user-level, fund-level and query-level characteristics. Investor characteristics are measured at the end of the previous month and include risk preference (*RiskPreference*), the natural logarithm of account value (*AccountValue*), portfolio diversification (*Diversification*), the proportion of equity funds in the investor's portfolio (*EquityFundRatio*), previous-month portfolio return (*L1mUserRet*), volatility (*L1mUserVolatility*), and turnover (*L1mUserTurnover*). Fund characteristics include previous-day return (*L1dFundRet*), previous-month return (*L1mFundRet*), and volatility (*L1mFundVolatility*), log size of the assets under management (*FundSize*), the log of fund age (*FundAge*), expenses (*Expenses*), and previous-quarter fund net flow (*L1qFundNetFlow*). Query-level characteristics includes user's trading intention. Standard errors are clustered at day level. Reported values in brackets are t-statistics. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Dep.var	*DUserTrade*		
Variables	(1)	(2)	(3)
AIRecTrade	0.163***	0.026***	0.024***
	[20.20]	[3.81]	[4.01]
RiskPreference		0.005***	0.005***
		[3.91]	[4.26]
AccountValue		-0.000	-0.000
		[-1.03]	[-0.26]
Diversification		-0.006	-0.005
		[-1.33]	[-1.03]
EquityFundRatio		-0.029***	-0.030***
		[-5.91]	[-7.35]
L1mUserRet(%)		-0.001*	-0.000
		[-1.86]	[-0.61]
L1mUserVolatility(%)		-0.007**	-0.005**
		[-2.14]	[-2.11]
L1mUserTurnover(%)		0.000*	0.000*
		[1.76]	[1.84]
L1dFundRet(%)		-0.001	0.000
		[-0.95]	[0.15]

L1mFundRet(%)		0.000	0.000
		[0.46]	[1.61]
L1mFundVolatility(%)		-0.002	-0.003
		[-0.68]	[-1.47]
FundSize		0.004***	0.004***
		[6.50]	[6.65]
FundAge		-0.014***	-0.014***
		[-6.18]	[-6.36]
Expenses(%)		-0.001	-0.001
		[-0.25]	[-0.40]
L1qFundNetFlow		0.001***	0.001***
		[3.10]	[2.89]
Constant	-0.017***	0.043***	0.040***
	[-5.09]	[4.05]	[4.41]
Controls	NO	YES	YES
Day FE	NO	NO	YES
Adj R-squared	0.008	0.061	0.069
Obs.	253,740	253,740	253,724

Table 5. Determinants of Following AI Advice: Investor and AI Response Characteristics
This table reports the regression results on how investor and response characteristics are associated with the likelihood of following AI investment recommendations. We estimate the panel regressions as specified in Equation (4). The dependent variable, *DUserFollowAI*, equals one if investor follows the AI recommendation within 24 hours of the query: buying after a buy recommendation, selling after a sell recommendation, or not trading after a hold recommendation. Panel A focus on *DLowEdu* and *DAdviceSeeking*. *DLowEdu* is an indicator equal to one if the investor has less than a bachelor's degree. *DAdviceSeeking* equals one if the user's primary intent is to seek investment advice rather than other objectives (e.g., seeking information). Panel B focus on *AIStrength*, *AIReadability*, *AIRelevance and AISatisfaction* as the main independent variables. *AIStrength* measures strength of recommendation, ranging from 0 to 1. *AIReadability* is an LLM-generated measure of the linguistic readability of the AI response, with higher values indicating responses that are easier to understand. *AIRelevance* is an LLM-generated measure of how relevant the AI response is to the investor's query, with higher values indicating better alignment between the response and the question. *AISatisfaction* measures the investor's expressed satisfaction with the AI response, taking the value of 1 for satisfied, −1 for dissatisfied, and 0 if no satisfaction or dissatisfaction is expressed. Controls include lagged investor characteristics (risk preference, account value, diversification, equity fund ratio, portfolio return, volatility, and turnover), fund characteristics (returns, volatility, size, age, expenses, and net flows), and user trading intention. Day fixed effect is included in all regressions. Standard errors are clustered at day level. Reported values in brackets are t-statistics. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. User characteristics

Dep.var	*DUserFollowAI*		
Variables	(1)	(2)	(3)
DLowEdu	0.012***		0.012***
	[5.21]		[5.10]
DAdviceSeeking		0.043***	0.046***
		[12.75]	[11.91]
RiskPreference	-0.001	-0.001	-0.001
	[-1.42]	[-1.31]	[-1.40]
AccountValue	-0.001	0.000	-0.001*
	[-1.61]	[0.14]	[-1.76]
Diversification	-0.028***	-0.025***	-0.027***
	[-5.41]	[-6.18]	[-5.30]
EquityFundRatio	-0.008**	-0.008***	-0.010***
	[-2.36]	[-2.61]	[-2.81]
L1mUserRet(%)	-0.000	-0.000	-0.000
	[-0.26]	[-1.31]	[-0.16]
L1mUserVolatility(%)	-0.002	-0.003	-0.003
	[-1.12]	[-1.43]	[-1.26]
L1mUserTurnover(%)	-0.000***	-0.000***	-0.000***
	[-9.91]	[-11.46]	[-9.35]
L1dFundRet(%)	-0.001	-0.001	-0.000
	[-0.61]	[-0.76]	[-0.52]
L1mFundRet(%)	-0.000	-0.000	-0.000
	[-0.48]	[-0.34]	[-0.27]
L1mFundVolatility(%)	-0.011***	-0.010***	-0.011***
	[-5.99]	[-6.81]	[-6.07]
FundSize	0.001	0.000	0.001
	[1.10]	[0.56]	[0.84]
FundAge	-0.004**	-0.003*	-0.004**
	[-2.07]	[-1.94]	[-1.98]
Expenses(%)	0.005*	0.004*	0.003
	[1.93]	[1.79]	[1.41]
L1qFundNetFlow	-0.001	-0.000	-0.001*
	[-1.64]	[-0.78]	[-1.65]
Constant	0.104***	0.090***	0.092***
	[11.69]	[12.10]	[10.80]
Controls	YES	YES	YES
Day FE	YES	YES	YES
Adj. R-squared	0.558	0.568	0.559
Obs.	155,924	253,724	155,924

Panel B. AI response characteristics

Dep.var	*DUserFollowAI*				
Variables	(1)	(2)	(3)	(4)	(5)
AIStrength	0.041***				0.040***
	[6.89]				[6.79]
AIReadability		0.068***			0.044***
		[4.37]			[2.78]
AIRelevance			0.078***		0.060***
			[5.09]		[3.81]
AISatisfaction				0.012***	0.013***
				[2.93]	[2.97]
RiskPreference	-0.001	-0.001	-0.001	-0.001	-0.001
	[-1.37]	[-1.32]	[-1.32]	[-1.32]	[-1.38]
AccountValue	0.000	0.000	0.000	0.000	0.000
	[0.32]	[0.26]	[0.15]	[0.23]	[0.26]
Diversification	-0.026***	-0.025***	-0.025***	-0.025***	-0.026***
	[-6.39]	[-6.30]	[-6.28]	[-6.31]	[-6.36]
EquityFundRatio	-0.006**	-0.006*	-0.006**	-0.006**	-0.006*
	[-2.04]	[-1.84]	[-2.07]	[-2.03]	[-1.95]
L1mUserRet(%)	-0.000	-0.000	-0.000	-0.000	-0.000
	[-1.38]	[-1.44]	[-1.42]	[-1.42]	[-1.38]
L1mUserVolatility(%)	-0.003	-0.002	-0.003	-0.002	-0.003
	[-1.32]	[-1.29]	[-1.30]	[-1.29]	[-1.31]
L1mUserTurnover(%)	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***
	[-12.18]	[-12.18]	[-11.97]	[-12.05]	[-12.19]
L1dFundRet(%)	-0.001	-0.001	-0.001	-0.001	-0.001
	[-0.89]	[-0.89]	[-0.84]	[-0.84]	[-0.95]
L1mFundRet(%)	-0.000	-0.000	-0.000	-0.000	-0.000
	[-0.40]	[-0.57]	[-0.57]	[-0.53]	[-0.46]
L1mFundVolatility(%)	-0.011***	-0.010***	-0.010***	-0.010***	-0.010***
	[-6.93]	[-6.55]	[-6.70]	[-6.72]	[-6.80]
FundSize	0.000	0.000	0.000	0.000	0.000
	[0.92]	[0.66]	[0.63]	[0.94]	[0.50]
FundAge	-0.003*	-0.003*	-0.003**	-0.003**	-0.003*
	[-1.79]	[-1.89]	[-2.10]	[-2.04]	[-1.76]
Expenses(%)	0.004**	0.006***	0.005**	0.005**	0.004**
	[1.97]	[2.68]	[2.28]	[2.35]	[2.11]
L1qFundNetFlow	-0.000	-0.000	-0.000	-0.000	-0.000
	[-0.71]	[-0.79]	[-0.80]	[-0.79]	[-0.73]
Constant	0.089***	0.049***	0.035**	0.100***	0.004
	[11.17]	[3.22]	[2.34]	[12.85]	[0.22]
Controls	YES	YES	YES	YES	YES
Day FE	YES	YES	YES	YES	YES
Adj. R-squared	0.568	0.567	0.567	0.567	0.568
Obs.	253,724	253,724	253,724	253,724	253,724

Table 6. AI Recommendations and Fund Future Performance

This table examines whether AI-generated fund recommendations predict subsequent fund performance. We estimate the panel regressions as specified in Equation (5). The dependent variables are the next-day fund return (*F1dFundRet*) in column (1), the next-week fund return (*F1wFundRet*) in column (2), and the next-month fund return (*F1mFundRet*) in column (3). The main independent variable is *AIRecTrade*, which captures both the direction and strength of the AI-generated fund recommendation, ranging from -1 to 1. Fund characteristics include previous-day, previous-month and previous-3month return, previous-month return volatility (*L1mFundVolatility*), log size of the assets under management (*FundSize*), the log of fund age (*FundAge*), expenses (*Expenses*), and previous-quarter fund net flow (*L1qFundNetFlow*). All regressions control for question-level trading intent, including buy, sell, and hold intent. Day fixed effect is included in all regressions. Standard errors are double clustered at fund and day level. Coefficients marked with *, **, and *** are significant at the 10%, 5%, and 1% levels, respectively.

Dep.var	*F1dFundRet(%)*	*F1wFundRet(%)*	*F1mFundRet(%)*
Variables	(1)	(2)	(3)
AIRecTrade	0.036	0.109	-0.455
	[0.60]	[0.68]	[-1.37]
FundRet(%)	-0.026	-0.065	-0.201*
	[-0.65]	[-0.53]	[-1.83]
L1mFundRet(%)	-0.005	-0.026	0.051*
	[-0.57]	[-1.12]	[1.86]
L3mFundRet(%)	-0.001	-0.013	-0.048***
	[-0.19]	[-0.99]	[-3.02]
L1mFundVolatility(%)	0.079	0.335***	0.451*
	[1.47]	[2.87]	[1.79]
FundSize	-0.002	0.006	-0.019
	[-0.22]	[0.27]	[-0.32]
FundAge	0.025	0.041	-0.055
	[0.95]	[0.61]	[-0.34]
Expenses(%)	-0.004	0.011	0.495*
	[-0.08]	[0.09]	[1.93]
L1qFundNetFlow	-0.000	0.015	0.049
	[-0.20]	[1.27]	[1.33]
Constant	0.036	0.046	0.910*
	[0.35]	[0.21]	[1.91]
Controls	Yes	Yes	Yes
Day FE	Yes	Yes	Yes
Adj. R-squared	0.359	0.269	0.333
Observations	253,293	253,293	253,293

Table 7. Following AI Advice and Subsequent Portfolio Returns

This table reports the association between following AI investment recommendations and investors' subsequent portfolio return. We estimate the panel regressions as specified in Equation (6). The dependent variable is *Portfolio Return* measured over future investment horizons of one day, one week, and one month following each query. *Portfolio Return* is the cumulative portfolio return over the corresponding horizon, reported in percentage terms. The main independent variable, *DUserFollowAI*, is an indicator equal to one if the investor follows the AI recommendation within 24 hours after submitting the query. Specifically, *DUserFollowAI* equals one if the investor buys following a buy recommendation, sells following a sell recommendation, or does not trade following a hold recommendation. *DBuyFollowAI*, *DSellFollowAI*, and *DHoldFollowAI* decompose *DUserFollowAI* by recommendation type. For example, *DBuyFollowAI* equals one if the AI issues a buy recommendation and the investor buys the recommended fund within 24 hours. All specifications control for user-, fund-, and query-level characteristics. User-level controls, measured at the end of the previous month, include *RiskPreference*, *AccountValue* (the natural logarithm of account value), Diversification, *EquityFundRatio*, previous-month portfolio return (*L1mUserRet*), volatility (*L1mUserVolatility*), and turnover (*L1mUserTurnover*). Fund-level controls include previous-day return (*L1dFundRet*), previous-month return (*L1mFundRet*), previous-month volatility (*L1mFundVolatility*), *FundSize* (the logarithm of assets under management), *FundAge* (the logarithm of fund age), Expenses (%), and previous-quarter fund net flow (*L1qFundNetFlow*). Query-level controls include the investor's trading intention. All specifications include day fixed effects. Standard errors are clustered at day level. Reported values in brackets are t-statistics. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Dep.var	*Portfolio Return* (%)					
Time Period	Future 1 day		Future 1 week		Future 1 month	
Variables	(1)	(2)	(3)	(4)	(5)	(6)
DUserFollowAI	0.018***		0.073***		0.109***	
	[2.83]		[3.98]		[3.33]	
DBuyFollowAI		-0.001		-0.003		-0.159**
		[-0.11]		[-0.07]		[-2.37]
DSellFollowAI		0.006		-0.030		-0.105
		[0.18]		[-0.25]		[-0.50]
DHoldFollowAI		0.024***		0.095***		0.175***
		[3.12]		[4.36]		[4.40]
RiskPreference	0.001	0.001	-0.013	-0.013	-0.011	-0.010
	[0.50]	[0.52]	[-1.52]	[-1.50]	[-0.46]	[-0.43]
AccountValue	-0.001	-0.001	0.008	0.008	-0.006	-0.006
	[-0.17]	[-0.16]	[0.81]	[0.82]	[-0.53]	[-0.51]
Diversification	0.013	0.013	0.062	0.063	0.480***	0.484***
	[0.83]	[0.85]	[1.40]	[1.43]	[4.36]	[4.40]

EquityFundRatio	0.026	0.026	0.317***	0.316***	1.615***	1.615***
	[0.62]	[0.62]	[3.08]	[3.08]	[8.55]	[8.55]
L1mUserRet(%)	0.000	0.000	0.045***	0.045***	0.100***	0.100***
	[-0.01]	[-0.01]	[4.13]	[4.13]	[4.39]	[4.39]
L1mUserVolatility(%)	-0.025	-0.026	-0.050	-0.051	-0.492***	-0.494***
	[-0.59]	[-0.59]	[-0.48]	[-0.49]	[-3.87]	[-3.89]
L1mUserTurnover(%)	0.000	0.000	0.000	0.000	-0.001	-0.001
	[-1.19]	[-1.17]	[0.10]	[0.12]	[-1.53]	[-1.48]
L1dFundRet(%)	0.006	0.006	0.002	0.002	-0.021	-0.021
	[0.53]	[0.53]	[0.07]	[0.07]	[-0.59]	[-0.58]
L1mFundRet(%)	-0.003	-0.003	-0.014**	-0.014**	0.004	0.004
	[-1.04]	[-1.04]	[-2.23]	[-2.23]	[0.53]	[0.53]
L1mFundVolatility(%)	0.020*	0.020*	0.105***	0.106***	0.051	0.053
	[1.65]	[1.67]	[3.92]	[3.95]	[0.79]	[0.82]
FundSize	0.000	0.000	0.004	0.004	0.005	0.006
	[0.01]	[0.04]	[0.86]	[0.91]	[0.57]	[0.66]
FundAge	0.003	0.003	-0.009	-0.009	-0.056*	-0.057*
	[0.41]	[0.41]	[-0.49]	[-0.50]	[-1.91]	[-1.93]
Expenses(%)	0.002	0.002	0.033	0.033	0.189***	0.189***
	[0.19]	[0.20]	[1.50]	[1.52]	[3.88]	[3.88]
L1qFundNetFlow	-0.002	-0.002	0.000	0.000	0.009*	0.009*
	[-1.55]	[-1.54]	[0.07]	[0.09]	[1.85]	[1.88]
Constant	0.013	0.013	0.022	0.022	0.923***	0.928***
	[0.20]	[0.20]	[0.12]	[0.12]	[4.39]	[4.41]
Controls	YES	YES	YES	YES	YES	YES
Day FE	YES	YES	YES	YES	YES	YES
Adj. R-squared	0.481	0.481	0.415	0.415	0.443	0.443
Obs.	248,581	248,581	234,066	234,066	188,638	188,638

Table 8. Following AI Advice and Portfolio Risk, Sharpe Ratios, and Turnover

This table reports the association between following AI investment recommendations and investors' subsequent investment performance. We estimate the panel regressions as specified in Equation (6). The dependent variables are *Portfolio Volatility*, *Portfolio Sharpe Ratio*, and *Portfolio Turnover*, measured over future investment horizons of one week and one month following each query. *Portfolio Volatility(Vol)* is the standard deviation of daily portfolio returns over the horizon. *Portfolio Sharpe Ratio(SR)* is the ratio of portfolio return minus risk-free rate to the standard deviation of daily portfolio returns over the same horizon. *Portfolio Turnover(TO)* is measured as the cumulative trading amount divided by the maximum portfolio value over the corresponding horizon. *Vol* and *TO* are multiplied by 100 for ease of interpretation. The main independent variable in Panel A, *DUserFollowAI*, is an indicator equal to one if the investor follows the AI recommendation within 24 hours after submitting the query. Specifically, *DUserFollowAI* equals one if the investor buys following a buy recommendation, sells following a sell recommendation, or does not trade following a hold recommendation. In Panel B, *DBuyFollowAI*, *DSellFollowAI*, and *DHoldFollowAI* decompose *DUserFollowAI* by recommendation type. *DBuyFollowAI* equals one if the investor buys following a buy recommendation, *DSellFollowAI* equals one if the investor sells following a sell recommendation, and *DHoldFollowAI* equals one if the investor does not trade following a hold recommendation, all measured within 24 hours after the query. All specifications control for user-, fund-, and query-level characteristics. User-level controls, measured at the end of the previous month, include *RiskPreference*, *AccountValue* (the natural logarithm of account value), Diversification, *EquityFundRatio*, previous-month portfolio return (*L1mUserRet*), volatility (*L1mUserVolatility*), and turnover (*L1mUserTurnover*). Fund-level controls include previous-day return (*L1dFundRet*), previous-month return (*L1mFundRet*), previous-month volatility (*L1mFundVolatility*), *FundSize* (the logarithm of assets under management), *FundAge* (the logarithm of fund age), Expenses, and previous-quarter fund net flow (*L1qFundNetFlow*). Query-level controls include the investor's trading intention. All specifications include day fixed effect. Standard errors are clustered at day level. Reported values in brackets are t-statistics. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Overall AI recommendation following

Dep.var	Vol (%)		SR		TO (%)	
Time Period	Future 1 week	Future 1 month	Future 1 week	Future 1 month	Future 1 week	Future 1 month
Variables	(1)	(2)	(3)	(4)	(5)	(6)
DUserFollowAI	0.005	0.006	0.006**	0.002**	-10.460***	-13.595***
	[1.12]	[1.26]	[2.46]	[1.97]	[-28.09]	[-22.42]
RiskPreference	-0.003	-0.002	-0.000	0.000	-0.229	-0.379
	[-1.27]	[-0.66]	[-0.34]	[0.39]	[-1.51]	[-0.65]
AccountValue	-0.026***	-0.024***	0.005***	0.003***	-1.683***	-2.854***
	[-13.18]	[-14.98]	[3.56]	[6.53]	[-25.49]	[-12.79]

Diversification	-0.085***	-0.060***	-0.019	-0.011	-4.496***	-7.496***
	[-5.83]	[-4.41]	[-1.17]	[-1.33]	[-7.44]	[-3.81]
EquityFundRatio	0.217***	0.251***	-0.009	0.007	0.541	-1.125
	[9.89]	[15.82]	[-0.55]	[0.98]	[1.08]	[-0.72]
L1mUserRet(%)	0.012***	0.009***	0.005***	0.003***	-0.034	-0.160
	[3.98]	[5.52]	[4.11]	[4.91]	[-0.66]	[-1.06]
L1mUserVolatility(%)	0.433***	0.403***	-0.057***	-0.058***	1.254***	2.628*
	[24.90]	[24.11]	[-8.83]	[-15.05]	[3.49]	[1.95]
L1mUserTurnover(%)	0.000***	0.001***	-0.000	-0.000***	0.139***	0.529***
	[7.23]	[10.79]	[-0.87]	[-4.72]	[37.69]	[43.37]
L1dFundRet(%)	-0.008*	-0.002	0.002	0.000	0.161**	0.251
	[-1.96]	[-0.77]	[0.66]	[0.12]	[2.12]	[1.42]
L1mFundRet(%)	0.006***	0.004***	-0.001	0.000	0.116***	0.388***
	[7.20]	[6.35]	[-1.04]	[1.55]	[6.39]	[7.21]
L1mFundVolatility(%)	0.068***	0.075***	-0.003	-0.012***	0.138	2.710***
	[11.39]	[15.19]	[-1.02]	[-6.79]	[1.03]	[5.74]
FundSize	0.000	-0.001	0.000	0.001	-0.373***	-1.603***
	[0.43]	[-1.14]	[0.05]	[1.39]	[-7.30]	[-10.15]
FundAge	-0.002	-0.006**	-0.009***	-0.011***	-1.124***	-2.890***
	[-0.61]	[-2.02]	[-3.34]	[-8.26]	[-7.71]	[-6.60]
Expenses(%)	0.054***	0.067***	0.003	0.000	-0.060	-1.953**
	[10.23]	[13.91]	[0.76]	[0.08]	[-0.28]	[-2.52]
L1qFundNetFlow	0.000	-0.000	0.000	0.000	0.063**	0.337***
	[0.84]	[-0.52]	[0.50]	[1.38]	[2.31]	[4.61]
Constant	0.461***	0.474***	0.155***	0.148***	35.618***	82.075***
	[19.23]	[22.11]	[7.76]	[18.01]	[40.91]	[27.75]
Controls	YES	YES	YES	YES	YES	YES
Day FE	YES	YES	YES	YES	YES	YES
Adj. R-squared	0.544	0.519	0.400	0.401	0.169	0.303
Obs.	233,589	188,347	229,697	187,834	235,347	189,365

Panel B. AI recommendation following by recommendation type

Variables	Vol (%)		SR		TO (%)	
Time Period	Future 1 week	Future 1 month	Future 1 week	Future 1 month	Future 1 week	Future 1 month
Variables	(1)	(2)	(3)	(4)	(5)	(6)
DBuyFollowAI	-0.028***	-0.032***	-0.002	-0.004	9.994***	12.085***
	[-3.37]	[-3.75]	[-0.38]	[-1.48]	[17.85]	[8.14]
DSellFollowAI	-0.046**	-0.007	0.004	-0.007	16.982***	19.101***
	[-2.40]	[-0.32]	[0.28]	[-0.78]	[9.69]	[5.99]
DHoldFollowAI	0.015***	0.015***	0.009***	0.004***	-16.391***	-20.147***
	[2.67]	[2.74]	[2.85]	[2.88]	[-42.93]	[-32.99]
RiskPreference	-0.003	-0.002	-0.000	0.000	-0.281*	-0.456
	[-1.23]	[-0.63]	[-0.32]	[0.42]	[-1.90]	[-0.79]
AccountValue	-0.026***	-0.024***	0.005***	0.003***	-1.699***	-2.873***
	[-13.17]	[-14.95]	[3.57]	[6.54]	[-26.27]	[-12.97]
Diversification	-0.085***	-0.060***	-0.019	-0.011	-4.909***	-7.959***
	[-5.79]	[-4.36]	[-1.16]	[-1.32]	[-8.24]	[-4.07]
EquityFundRatio	0.217***	0.251***	-0.009	0.007	0.583	-1.086
	[9.89]	[15.81]	[-0.55]	[0.98]	[1.17]	[-0.70]
L1mUserRet(%)	0.012***	0.009***	0.005***	0.003***	-0.035	-0.159
	[3.98]	[5.52]	[4.11]	[4.91]	[-0.70]	[-1.06]
L1mUserVolatility(%)	0.433***	0.403***	-0.057***	-0.058***	1.436***	2.837**
	[24.88]	[24.09]	[-8.84]	[-15.07]	[4.05]	[2.13]
L1mUserTurnover(%)	0.000***	0.001***	-0.000	-0.000***	0.137***	0.526***
	[7.28]	[10.86]	[-0.85]	[-4.68]	[37.57]	[43.54]
L1dFundRet(%)	-0.008*	-0.002	0.002	0.000	0.141*	0.223
	[-1.95]	[-0.76]	[0.66]	[0.13]	[1.92]	[1.29]
L1mFundRet(%)	0.006***	0.004***	-0.001	0.000	0.119***	0.387***
	[7.19]	[6.35]	[-1.04]	[1.56]	[6.87]	[7.20]
L1mFundVolatility(%)	0.068***	0.076***	-0.003	-0.012***	-0.061	2.485***
	[11.43]	[15.26]	[-0.99]	[-6.76]	[-0.46]	[5.39]
FundSize	0.001	-0.001	0.000	0.001	-0.440***	-1.681***

	[0.53]	[-1.03]	[0.08]	[1.44]	[-8.81]	[-10.68]
FundAge	-0.002	-0.006**	-0.009***	-0.011***	-1.056***	-2.838***
	[-0.64]	[-2.04]	[-3.36]	[-8.27]	[-7.45]	[-6.52]
Expenses(%)	0.054***	0.067***	0.003	0.000	-0.154	-1.961***
	[10.28]	[13.93]	[0.77]	[0.08]	[-0.77]	[-2.62]
L1qFundNetFlow	0.001	-0.000	0.000	0.000	0.046*	0.322***
	[0.88]	[-0.47]	[0.52]	[1.40]	[1.69]	[4.40]
Constant	0.461***	0.475***	0.155***	0.148***	35.498***	81.633***
	[19.25]	[22.13]	[7.75]	[18.01]	[41.89]	[28.06]
Controls	YES	YES	YES	YES	YES	YES
Day FE	YES	YES	YES	YES	YES	YES
Adj. R-squared	0.544	0.519	0.400	0.401	0.183	0.307
Obs.	233,589	188,347	229,697	187,834	235,347	189,365

Table 9. Response Quality and Further Engagement

This table reports the regression results on the determinants of investors' subsequent engagement with AI. We estimate the panel regressions as specified in Equation (7). The dependent variable, *FutureUse*, is defined as the natural logarithm of one plus the number of additional queries submitted by investor within one month following query. The main independent variables capture investor characteristics and query characteristics. *DLowEdu* is an indicator equal to one if the investor has less than a bachelor's degree. *DAdviceSeeking* equals one if the investor's primary intent is to seek investment advice rather than other objectives (e.g., seeking information). *AIStrength* measures the strength of the AI recommendation, ranging from 0 to 1. *AIReadability* is an LLM-generated measure of the linguistic readability of the AI response, with higher values indicating responses that are easier to understand. *AIRelevance* is an LLM-generated measure of how relevant the AI response is to the investor's query, with higher values indicating better alignment between the response and the question. *AISatisfaction* measures the investor's expressed satisfaction with the AI response, taking the value of 1 for satisfied, −1 for dissatisfied, and 0 if no satisfaction or dissatisfaction is expressed. All specifications control for user- and query-level characteristics. User-level controls, measured at the end of the previous month, include *RiskPreference*, *AccountValue* (the natural logarithm of account value), Diversification, *EquityFundRatio*, previous-month portfolio return (*L1mUserRet*), volatility (*L1mUserVolatility*), and turnover (*L1mUserTurnover*). Query-level controls include the investor's trading intention. Day fixed effects are included in all regressions. Standard errors are clustered at the day level. Reported values in brackets are t-statistics. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Dep.var	*FutureUse*						
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)
DLowEdu	0.335***						0.329***
	[5.44]						[5.49]
DAdviceSeeking		0.070***					0.043*
		[3.00]					[1.72]
AIStrength			0.038*				0.021
			[1.80]				[0.88]
AIReadability				-0.144			-0.163
				[-1.31]			[-1.29]
AIRelevance					0.152***		0.101**
					[4.14]		[2.41]
AISatisfaction						0.432***	0.429***
						[5.60]	[4.66]

RiskPreference	0.131***	0.102***	0.102***	0.102***	0.102***	0.101***	0.130***
	[4.78]	[4.90]	[4.89]	[4.89]	[4.90]	[4.99]	[4.90]
AccountValue	0.014**	0.018***	0.018***	0.018***	0.018***	0.018***	0.013**
	[2.35]	[3.44]	[3.47]	[3.46]	[3.45]	[3.64]	[2.31]
Diversification	0.601***	0.585***	0.585***	0.585***	0.586***	0.576***	0.593***
	[6.46]	[7.82]	[7.83]	[7.78]	[7.83]	[7.87]	[6.48]
EquityFundRatio	0.099	0.122**	0.126**	0.124**	0.125**	0.124**	0.091
	[1.51]	[2.34]	[2.38]	[2.35]	[2.37]	[2.41]	[1.45]
L1mUserRet(%)	0.001	0.003	0.003	0.003	0.003	0.003	0.001
	[0.34]	[0.70]	[0.70]	[0.70]	[0.69]	[0.72]	[0.35]
L1mUserVolatility(%)	-0.110***	-0.121***	-0.122***	-0.122***	-0.121***	-0.120***	-0.107***
	[-4.08]	[-4.22]	[-4.22]	[-4.26]	[-4.21]	[-4.26]	[-4.05]
L1mUserTurnover(%)	0.003***	0.003***	0.003***	0.003***	0.003***	0.003***	0.003***
	[4.93]	[8.46]	[8.45]	[8.45]	[8.46]	[8.76]	[5.15]
Constant	0.212*	0.449***	0.450***	0.570***	0.337***	0.445***	0.244
	[1.80]	[6.44]	[6.39]	[5.08]	[4.31]	[6.53]	[1.48]
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	0.131	0.121	0.121	0.121	0.121	0.131	0.141
Obs.	242,140	384,957	384,957	384,957	384,957	384,957	242,140

Appendix Table A1. Variable Definition

Variable	Definition
AI adoption and queries	
DAIUser	Indicator variable equal to one if investor *i* uses the AI investment assistant
NQuery	Natural logarithm of one plus the total number of AI queries submitted by
NInfoQuery	Natural logarithm of one plus the number of information-seeking AI queries
NAdviceQuery	Natural logarithm of one plus the number of advice-seeking AI queries
Investor characteristics	
Age	Investor age in years, measured as of October 2025.
Male	Indicator variable equal to one if the investor is male, and zero otherwise.
FamilyIncome	Categorical variable measuring the investor's annual household income
DLowEdu	Indicator variable equal to one if the investor has an educational attainment
DTier1City	Indicator variable equal to one if the investor resides in a Tier-1 city in
AccountValue	Natural logarithm of the investor's total financial assets CNY held in the
Diversification	Measure of portfolio diversification, calculated as one minus the Herfindahl
RiskPreference	Investor risk tolerance based on the platform's risk assessment survey,
AccountAge	Natural logarithm of the number of years since the investor opened the
DAutoinvestUser	Indicator variable equal to one if the investor uses the platform's automatic
InvestmentExp	Categorical variable measuring the investor's investment experience,
EquityFundRatio	Proportion of equity funds in the investor's fund portfolio.
UserRet(%)	Cumulative user portfolio return over the horizon.
UserVolatility(%)	Standard deviation of daily portfolio returns over the corresponding
Sharpe	The ratio of portfolio return minus risk-free rate to the standard deviation
UserTurnover(%)	Cumulative trading amount divided by the maximum portfolio value over
AI recommendations and subsequent trading	
AIRecTrade	Signed measure capturing the direction and strength of the AI-generated
DUserTrade	Trading-direction variable measured within 24 hours after the query, equal
DUserFollowAI	Indicator variable equal to one if the investor follows the AI
DAdviceSeeking	Indicator variable equal to one if the user's query intent is to seek
AI response characteristics	
AIStrength	Absolute value of LLM-generated measure of the AI recommendation tone.
AIReadability	LLM-generated measure of how easily a retail investor can understand the
AIRelevance	LLM-generated measure of how closely the AI response addresses the
AISatisfaction	Measure of the investor's expressed satisfaction in the first follow-up query
Fund characteristics	
FundSize	Natural logarithm of the fund's assets under management.
FundAge	Natural logarithm of one plus fund age (year).
Expenses(%)	Fund expenses, expressed as a percentage.
FundNetFlow	Fund net flow over the previous quarter.
L1dFundRet(%)	Fund return on the previous trading day.
L1mFundRet(%)	Fund return over the previous month.
L1mFundVolatility(%)	Volatility of daily fund returns over the previous month
L1qFundNetFlow	Fund net flow over the previous quarter.

Internet Appendix for

"Does Your AI Move Your Money?

Adoption and Impact of a Generative AI Investment Assistant"

August 2026

This internet appendix includes the following tables used in the paper.

Table IA1. Correlation Matrix: User Characteristics
Table IA2. Query Classification: User Intent
Table IA3. Examples of AI Queries and Answers
Table IA4. AI Recommendation and Fund Future Performance
Table IA5. Response Quality and Investor Satisfaction

Table IA1. Correlation Matrix: User Characteristics

This table reports the pairwise correlations among the main user characteristics. The sample period spans from July 2024 to October 2025 and consists of 400,000 randomly selected active individual investors on the Ant platform who executed at least one fund transaction or held at least RMB 100 in funds during the sample period. *Age* is investor age; *Male* indicates investor gender; *DLowEdu* indicates lower educational attainment; *FamilyIncome* is household income; *DTier1City* indicates Tier-1 city residence; *AccountValue* is total financial assets; *Diversification* measures portfolio diversification; *RiskPreference* measures risk tolerance; *AccountAge* is the number of years since account opening; *DAutoinvestUser* indicates auto-invest usage; and *InvestmentExp* measures investment experience. The reported values are Pearson correlation coefficients. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. Detailed variable definitions are provided in Appendix Table A1.

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
(1)*Age*	1										
(2)*Male*	0.009***	1									
(3)*DLowEdu*	0.296***	-0.049***	1								
(4)*FamilyIncome*	0.059***	-0.056***	-0.118***	1							
(5)*DTier1City*	-0.079***	0.027***	-0.147***	0.105***	1						
(6)*AccountValue*	0.080***	0.129***	-0.205***	0.199***	0.120***	1					
(7)*Diversification*	-0.035***	0.016***	-0.102***	0.076***	0.053***	0.211***	1				
(8)*RiskPreference*	-0.145***	-0.040***	-0.182***	0.326***	0.105***	0.300***	0.161***	1			
(9)*AccountAge*	0.142***	0.085***	-0.141***	0.161***	0.076***	0.293***	0.069***	0.287***	1		
(10)*DAutoinvestUser*	-0.028***	0.009***	-0.111***	0.084***	0.059***	0.230***	0.228***	0.165***	0.064***	1	
(11)*InvestmentExp*	-0.015***	0.016***	-0.108***	0.164***	0.068***	0.247***	0.093***	0.413***	0.348***	0.113***	1

Table IA2. Query Classification: User Intent

This table presents the classification of user queries by intent. Each query is assigned to a primary category (Category) and a more granular secondary category (Secondary Category). The secondary categories are based on the query-intent labels provided by Ant Group. We further aggregate these secondary categories into three primary categories: Information, which captures queries seeking information, interpretation, or other non-advisory assistance; Advice, which captures queries explicitly seeking investment advice; and Other, which captures casual chat, greetings, weather-related queries, and other queries unrelated to investment information or advice.

Category	Secondary Category
Information	Product information inquiry | Basic information (funds, fixed-income products, etc.)
Information	Product information inquiry | Funds | Fund comparison
Information	Product information inquiry | Funds | Fund managers
Information	Product information inquiry | Funds | Product evaluation
Information	Product information inquiry | Funds | Historical returns
Information	Product information inquiry | Funds | Return prediction
Information	Product information inquiry | Stocks | Corporate executives
Information	Product information inquiry | Stocks | Stock comparison
Information	Product information inquiry | Stocks | Stock indicators
Information	Product information inquiry | Stocks | Stock evaluation and opinions
Information	Product information inquiry | Stocks | Price indicators
Information	Product information inquiry | Stocks | Financial reports and events
Information	Product information inquiry | Other
Information	Product information inquiry | Related news and information
Information	Financial education / financial and non-financial terminology
Information	Market interpretation | Alternative investments
Information	Market interpretation | Cross-industry comparison
Information	Market interpretation | Objective indicators
Information	Market interpretation | Events and news
Information	Market interpretation | Future outlook
Information	Market interpretation | Market movement interpretation
Information	News and information | Other
Information	News and information | Future impact
Information	Market interpretation | Market index movement and interpretation
Information	Portfolio management | Performance analysis | Single product
Information	Portfolio management | Performance analysis | Overall portfolio
Information	Portfolio management | Portfolio information (watch-list, bills, holdings)
Information	Portfolio management | Financial score
Information	Portfolio management | Risk assessment
Information	Platform services | Product navigation
Information	Platform events and promotions
Information	Find a financial advisor
Information	Contact human customer service
Information	Customer service inquiries

Information	AI financial advisor
Information	Non-platform products
Information	About Zhixiaobao (platform assistant)
Information	Ant Group related questions
Advice	Product information inquiry | Funds | Buy/sell advice
Advice	News and information | Trading advice
Advice	Market interpretation | Trading advice
Advice	Portfolio management | Goal setting
Advice	Product selection | Other recommendations
Advice	Product selection | Fund manager recommendation
Advice	Product selection | Timing / scenario-based selection
Advice	Product selection | Conditional fund screening
Advice	Product selection | Conditional stock screening
Advice	Product selection | Industry recommendation
Advice	Product selection | Unspecified fund screening
Advice	Product selection | Unspecified stock screening
Advice	Portfolio allocation | Three-bucket framework
Advice	Portfolio allocation | Savings and protection funds (second bucket)
Advice	Portfolio allocation | Daily expenses funds (first bucket)
Advice	Portfolio allocation | Future growth funds (third bucket)
Advice	Portfolio allocation | Other
Other	Casual chat | Greetings
Other	Casual chat | Weather
Other	Casual chat
Other	Other

Table IA3. Examples of AI Queries and Answers

This table presents examples of answers of financial AI assistant to information-seeking and advice-seeking queries. The examples illustrate the general structure of the AI responses, which typically combine emotional support, factual information and analysis, and investment advice. Italicized text indicates emotional support or empathy, normal text presents factual information and analysis, and underlined text indicates investment advice.

Panel A. Information-seeking query and answer

Question: Help me understand the different sectors of the A-share market

Answer:

"I understand that you would like to identify investment opportunities across different A-share sectors under the current market environment. Overall, the A-share market is exhibiting a structurally differentiated pattern, with technology and growth sectors remaining active, although short-term volatility risks should be taken into account.

I. Overall Market Performance and Fund Flows

1. Major indices moved higher amid intensified competition among investors. As of the close on August 29, 2025, the Shanghai Composite Index stood at 3,857.93, up 0.37%. On that day, major-fund inflows amounted to RMB 510.602 billion, while outflows totaled RMB 543.775 billion, resulting in a net outflow of RMB 33.173 billion. Overall, the market remains in a period of consolidation and adjustment, with investors displaying a relatively cautious attitude.

…

V. Summary and Outlook

The A-share market is undergoing an important transition from quantitative expansion toward qualitative improvement, and the shift in asset-pricing logic from traditional industries toward technological innovation has become an irreversible trend. I suggest remaining patient, focusing on policy-supported technology and consumption sectors, and maintaining a diversified portfolio rather than putting all your eggs in one basket. In terms of specific investment strategies, you may consider participating in the market through regular investment plans to reduce market-timing risk."

In Chinese:

Question: 帮我解读 A 股板块

Answers:

"*我理解你想在当前市场环境下把握 A 股板块的投资机会。*综合来看， A 股市场呈现结构性行情，科技成长板块表现活跃，但需注意短期波动风险。

一、市场整体表现与资金动向

1.指数震荡上行，资金博弈加剧：截至 2025 年 8 月 29 日收盘，上证指数报 3857.93 点，上涨 0.37%。当日主力资金流入 5106.02 亿元，流出 5437.75 亿元，净流出 331.73 亿元。从趋势看，市场处于震荡调整阶段，资金呈现谨慎态度。

…

五、总结与展望

A 股市场正经历从量变到质变的重要阶段，资产定价逻辑从传统行业向科技创新转型已成为不可逆转的趋势。建议你保持耐心，关注政策支持的科技、消费等板块，同时注意分散配置，避免把鸡蛋放在一个篮子里。具体操作上，可以考虑通过定投方式参与市场，降低择时风险。”

Panel B. Advice-seeking query and answer

Question: Given the current market conditions, should I reduce my position in the XXX Fund?

Answer:

I understand your concerns about the current market volatility. However, based on your current holdings and our market analysis, the XXX Fund still offers allocation value. I suggest remaining patient and not rushing to reduce your position.

I. Market Environment and Product Positioning

The advantages of the dividend low-volatility strategy remain evident. As of August 26, 2025, the dividend low-volatility sector had a price-to-earnings ratio of 9.40.

…

Overall, the fund still offers value as a medium- to long-term allocation instrument. I suggest remaining patient and continuing to monitor market developments.

In Chinese:

Question: 这行情，我要减仓 XXX 基金吗？

Answers:

“*我理解你对当前市场波动的担忧，*但根据你的持仓情况和市场分析，XXX 基金仍具备配置价值，建议你保持耐心，不必急于减仓。

一、市场环境与产品定位

红利低波动策略优势凸显：截至 2025 年 08 月 26 日，红利低波动板块当前 PE 市盈率为 9.40

…

综合来看，该基金作为中长期配置工具仍具价值，建议你保持耐心，关注市场变化。”

Table IA4. AI Recommendation and Fund Future Performance

This table reports the results from the Fama-MacBeth (1973) regressions of future fund performance on the AI recommendations. The dependent variables are the next-day, next-week and next-month fund returns. The main independent variable, *BuyRatio*, is defined as the number of buy recommendations scaled by the total number of buy, sell, and hold recommendations for a given fund-day. Fund characteristics includes previous-day return (*L1dFundRet*), previous-month return (*L1mFundRet*), and volatility (*L1mFundVolatility*), log size of the assets under management (*FundSize*), the log of fund age (*FundAge*), expenses (*Expenses*), and previous-quarter fund net flow (*L1qFundNetFlow*). ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Dep.var	*F1dFundRet*	*F1wFundRet*	*F1mFundRet*
Variables	(1)	(2)	(3)
BuyRatio	-0.003	0.00006	-0.072
	[-0.36]	[0.17]	[-0.72]
Intercept	0.026*	0.0018***	0.825***
	[1.67]	[2.78]	[6.26]
Controls	Yes	Yes	Yes
Adj. R-squared	0.383	0.404	0.396

Table IA5. Response Quality and Investor Satisfaction

This table reports the relation between AI response quality and investor follow-up satisfaction. The dependent variable is *AISatisfaction*, which indicates whether investor expresses a positive verbal response in the first follow-up query submitted within 15 minutes after receiving AI response. Follow-up satisfaction is identified based on the content of this first follow-up query. If the user's follow-up response cannot be reliably classified as positive or negative, the observation is labeled as Unknown. The key independent variables are response-quality measures, including *Relevance* and *Readability*. *Relevance* measures how well the response addresses the user's question, and *Readability* measures how easy the response is for a retail investor to understand. Columns (1) to (2) examine each response-quality measure separately, and Column (3) includes all response-quality measures jointly. All regressions control for question-level trading propensity. Day fixed effects are included in all regressions. Standard errors are clustered at day level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Dep.var	*AISatisfaction*		
Variables	(1)	(2)	(3)
Relevance	0.024		0.013
	[1.15]		[0.62]
Readability		0.133***	0.149***
		[3.04]	[3.41]
Constant	0.311***	0.226***	0.187***
	[16.58]	[7.43]	[5.50]
Controls	Yes	Yes	Yes
Day FE	Yes	Yes	Yes
R-squared	0.286	0.287	0.287
Obs.	63,581	63,581	63,581