Review of Futures Markets
ISSN 1933-7116 (Print)
fra.st/rfm

Volume 34 · Issue 2 · 2026
Pages 1707–1718

Impact of AI-Based Investment Recommendations on Retail Investor Psychology Markets

Vazeer Jan Begum[1], Dr. Urvi Amin[2], Swati Yadav[3], Shalini Singh[4], Dr. Anchal Saini[5], Dr. S. Darshan[6], Jigar Rupani[7], Dr. Sandeep Dongre[8]

[1] Professor, Dubai Business School, University of Dubai

[2] Shri Jairambhai Patel Institute of Business Management and Computer Applications, (National Institute of Cooperative Management-NICM), Gujarat

[3] Assistant Professor, Graduate School of Business, Tula's Institute, Dehradun

[4] Assistant Professor, School of Commerce, Finance & Accountancy, CHRIST (Deemed to be University), India

[5] Assistant Professor, Department of Commerce, CHRIST (Deemed to be University), Delhi NCR, India

[6] Assistant Professor, School of Economics and Business, RV University, Bengaluru, Karnataka

[7] Assistant Professor, School of Business and Management, CHRIST UNIVERSITY

[8] Professor of Practice, Symbiosis Institute of Business Management (SIBM) Nagpur, constituent of Symbiosis International (Deemed University), Nagpur- 440008 Maharashtra India

Received: 18/09/2026 | **Accepted:** 21/10/2026 | **Published:** 28/10/2026

Abstract

Within roughly five years, AI investment-recommendation systems — from rule-based robo-advisors to large-language-model chatbots such as ChatGPT, Gemini and Copilot — have moved from the margins to the mainstream of retail investing. This paper synthesises evidence on how receiving such recommendations affects the psychology of non-professional investors. It draws on a behavioural-science experiment of 7,771 Canadians conducted by the Ontario Securities Commission and the Behavioural Insights Team; structural-equation-modelling studies of 461 and 409 retail investors in India; a 938-respondent U.S. investor survey; and a Federal Reserve working paper on the financial-stability implications of generative AI. We report four findings. First, adoption among retail investors is rapid and widespread: roughly 62% of U.S. retail investors reported using AI tools in 2025–2026, and about 48% of Australian retail investors reported the same. Second, trust is not binary: most users trust AI only "somewhat" and continue to verify its outputs against other sources, so nuanced trust — rather than wholesale delegation — drives adoption. Third, robo-advisors can reduce overconfidence and loss-aversion biases when they are perceived as personalised, interactive and transparent; herding, however, is not consistently diminished. Fourth, the same systems introduce new psychological risks: automation bias, FOMO-driven over-reliance, and the possibility that correlated AI outputs amplify herd-like behaviour in markets. It is been concluded that AI recommendations behave less like a substitute for human judgement and more like a powerful new heuristic, with their net psychological impact depending on system design, user financial literacy and the regulatory framework.

Keywords: artificial intelligence; retail investors; robo-advisors; behavioural finance; algorithm aversion; overconfidence bias; investor psychology

How to cite: Begum et al. (2026). Impact of AI-Based Investment Recommendations on Retail Investor Psychology Markets. Review of Futures Markets, 34(2), 1707–1718.
© 2026 The Author(s). Published by Review of Futures Markets.

1. Introduction

For the better part of the twentieth century, the retail investor was treated as a trivial sideshow in capital markets — a player whose behaviour could only be analysed on an aggregate basis and who had little impact on price dynamics from his own psychology. That image has become an artifact of the past five years. The past few years have shown us that retail flows can even be enough to move prices in ways then-current classical models (like the ones Nadig describes) did not anticipate: for instance, during the COVID-19 pandemic, with zero-commission brokerage apps flooding order books and during January 2021's meme-stock episode. The transformation accelerated with the arrival of generative artificial intelligence in late 2022. In a survey conducted in March 2026, most U.S. retail investors (62%) stated that they were using AI tools to aid their investment decisions and more than half (53.5%), reported investing with the help of conversational chatbot style across ChatGPT as one example for investment research (Monteiro, 2026).

Such swift spread has been achieved, but without — or perhaps because of — any developed regulatory mechanism. Currently, no commercially available AI tool is approved under the Markets in Financial Instruments Directive (MiFID) to provide investment advice in the EU even though more than one-in-five EU retail investors currently use AI tools or services when adjusting portfolio decisions (European Securities and Markets Authority, 2025; eToro, 2025). The Canadian story is no different: while AI may not be allowed to give advice without the supervision of a human, one Ontario Securities Commission (OSC) analysis found that Canadians were too inharmonious about whether a suggestion neural-net by either a human or an AI from a combined human–AI process (Ontario Securities Commission [OSC] & Behavioural Insights Team [BIT], 2024).

This paper does not address whether any AI-based investment recommendations are actually accurate, an econometric question. What we seek instead is the following: how does a suggestion of these recommendations alter the psychology of retail investors? Three sub-questions structure the analysis:

RQ1How and when do ordinary investors learn to trust an AI advisor—and how does that trust compare to what they'd give a human?

RQ2. How do AI recommendations mitigate or amplify specific classical cognitive biases — overconfidence, loss aversion, herding, anchoring and availability?

RQ3. What new psychological phenomena emerge from the human–AI advisory relationship, and what are the implications for market stability?

The paper is integrative in its contribution. Instead of a new main data-set, it distils evidence from twelve recent empirical studies — including three with sample sizes exceeding 900 and one with 7,771 respondents — to chart the psychological environment for AI-driven retail investing. The rest of the paper is organised as follows. The next section (2) covers the literature on behavioural-finance and human–AI-interaction. In Section 3 we present the conceptional and methodological framework. Section 4 presents findings on adoption, trust and emergent risks related to bias mitigation. Section 5 outlines the theoretical and policy implications. Section 6 outlines limitations and an agenda for future research.

2. Literature Review

2.1 The Behavioural-Finance Foundation

For four decades, the behavioural-finance literature has documented systematic departures from the rational-agent model of investment decision-making. Kahneman and Tversky’s (1979) prospect theory established that investors weigh losses roughly twice as heavily as equivalent gains — a tendency known as loss aversion that distorts portfolio choice and trading. Barber and Odean (2001) showed that excessive trading is associated with overconfidence — particularly among men and younger investors — and leads to lower net returns. Herding, in which investors copy others rather than act on private information, can inflate asset-price bubbles (Shiller, 2000). The earlier heuristics-and-biases programme of Tversky and Kahneman (1974) had already established anchoring (over-reliance on an initial reference point) and the availability heuristic (overweighting easily recalled information) as pervasive features of human judgement under uncertainty.

These biases do not operate in a vacuum. In a review of robo-advisor research, Bhatia, Chandani and Chhateja (2020) note how fear of missing out (FOMO), social proof and digital echo chambers can magnify traditional biases. Survey evidence on Indian retail investors reports that a majority admit to following the

herd at least occasionally and that many experience FOMO when social-media discussion of a particular stock intensifies (see Section 4; source to be verified).

2.2 The Emergence of AI-Based Investment Recommendations

The available evidence groups AI applications in retail investing into three broad categories (OSC & BIT, 2024): decision support (systems that recommend allocations but leave execution to the user); automation (the traditional robo-advisor model, in which systems autonomously execute and rebalance portfolios); and adversarial uses, including AI-enabled fraud. Decision support has grown rapidly since the public release of ChatGPT in November 2022; by 2026, 53.5% of U.S. retail investors had used a general-purpose chatbot for investment research (Monteiro, 2026).

One of the more mature and narrower types of AI-based recommendation — robo-advisor — has increased too. Another more recent article by D'Acunto and Rossi (2023), is that robo-advisors made it possible for many more households to participate in a diversified-market portfolio. AI is no longer just a feature but a differentiator: Investing platforms such as. Generative AI is built into com (with WarrenAI and ProPicks AI), BridgeWise (track over 50,000 assets) as well as Robindood and eToro have integrated it (Investing.) com, 2026; Diamant, 2025).

2.3 Algorithm Aversion versus Algorithm Appreciation

How retail investors behave when AI advice is available to them is as yet a matter of conjecture, and the results are not uniform. In an experimental study, Dietvorst, Simmons and Massey (2015) discovered that one could train on a data set up until October 2023 when encountering forecasts made by algorithms to prefer the human even if — objectively speaking — the forecast performed worse; in other words they labelled what they observed as “algorithm aversion.” In contrast, the evidence of “algorithm appreciation”, where people made comparisons between advice that was identical apart from its labelled source and relied more on advice presented as algorithmic than equivalent advice given a human label: Logg et al.

Later research has started finding a compromise between these positions. In MIS Quarterly, Jussupow, Benbasat and Heinzl (2024) argue that aversion and appreciation are not opposing dispositions but rather context-dependent responses observed in the same individual. Castelo, Bos, and Lehmann (2019) demonstrate that the trade-off is conditional on task type: people prefer algorithms for analytical, quantitative tasks but develop an aversion to algorithms in cases where they think of a task as subjective. Expanding on this, Dietvorst, Simmons and Massey (2018) showed that granting users a small amount of control — the chance to adjust an algorithmic forecast by at most 10% — led to huge increases in their willingness to use it.

2.4 Robo-Advisors as Bias Mediators

A small but growing literature treats AI-backed robo-advisors as moderators or mediators of classic biases, rather than as mere delivery channels. Kulkarni, Patil and Pramod (2025) applied structural equation modelling to data from 461 retail investors in India and found that perceived personalisation, interactivity, autonomy and algorithm interpretability significantly reduced overconfidence and loss-aversion biases, whereas perceived structural assurance — the extent to which participants felt institutional safeguards were present — did not. In a study of New York Stock Exchange retail investors, Ahmad, van Keulen, Briassouli and Saad (2025) found that robo-advisor mediation attenuated several cognitive biases, including anchoring and overconfidence.

Adoption is itself biased, however. Hodge, Mendoza and Sinha (2021) and Loos, Previtero and Scheurle (2020) report that overconfident investors are disproportionately likely to adopt robo-advisors in the first place, which raises a selection problem: do robo-advisors reduce overconfidence, or do they simply attract already-overconfident investors willing to take systematic advice? In a study of 409 Indian investors, Verma, Schulze, Goswami and Upreti (2025) found that several barriers drive resistance to robo-advisor adoption — with inertia the strongest — and that a favourable attitude toward AI significantly weakened the effect of inertia, overconfidence bias and data-privacy concerns on that resistance, without eliminating it.

2.5 The Research Gap

This synthesis is motivated by three gaps. So far, most studies have focused on one or the other: robo-advisors, which are older and more narrowly built, or generative-AI chatbots. Few have brought both into the same frame. Second, adoption and trust have received

more attention than downstream psychological consequences such as automation bias, dependence and reduced accountability. Third, the financial-stability literature on AI herding (Hansen & Lee, 2025) has developed largely separately from the behavioural-finance literature on individual investor psychology. This paper brings these strands together into a single integrative account.

3. Conceptual Framework and Methodology

This paper uses a narrative integrative-review methodology rather than a systematic review with a registered search protocol. Its aim is not to test a single hypothesis on primary data, but to synthesise quantitative and qualitative findings from the most recent peer-reviewed and grey-literature sources directly concerned with retail-investor psychology under AI-mediated recommendation. No formal database search was conducted and no PRISMA-style flow diagram is presented; the source set reflects purposive sampling against the criteria below and therefore cannot claim completeness.

Sources were included if they (a) were published from January 2019 until April 2026; (b) concentrated on retail (non-professional), rather than institutional investors, and (c) provided quantitative evidence on at least one of the dependent variables of interest — trust, adoption, behavioural bias, adherence to advice or reported performance impact. The last set consists of twelve primary studies, the OSC/BIT (2024) experiment, Investing. In the Kulkarni et al. (quote) 2026 survey of 938 U.S. retail investors, (Monteiro. Verma et al. (2025) a SEM study of 461 Indian investors, (2025) a study with 409 Indian investors.

This conceptual structure consists of three layers. You are trained on (i) the source of a recommendation: human, AI or blended; and(ii) robustness features: personalisation, interpretability and accuracy. The second layer is the mediator layer: trust formation, perceived risk, perceived control and emotion. The response layer has the behavioural dimension to it: bias intensity, compliance as per requests made, concentration of trading activity and subjective performance. We map the results provided in Section 4 to one or more elements of this framework.

Due to the manner in which effect sizes cannot be pooled across studies that operationalise constructs like “trust” or “overconfidence,” this is done using an integrative rather than meta-analytic approach. They are included in the results a where standardised path coefficients have been reported in the literature but are not treated as fixed effect sizes, rather for illustrative purposes. Therefore, although the specific values are important, the discussion in Section 5 addresses direction and consistency of effects rather than exact quantitative numbers.

4. Findings

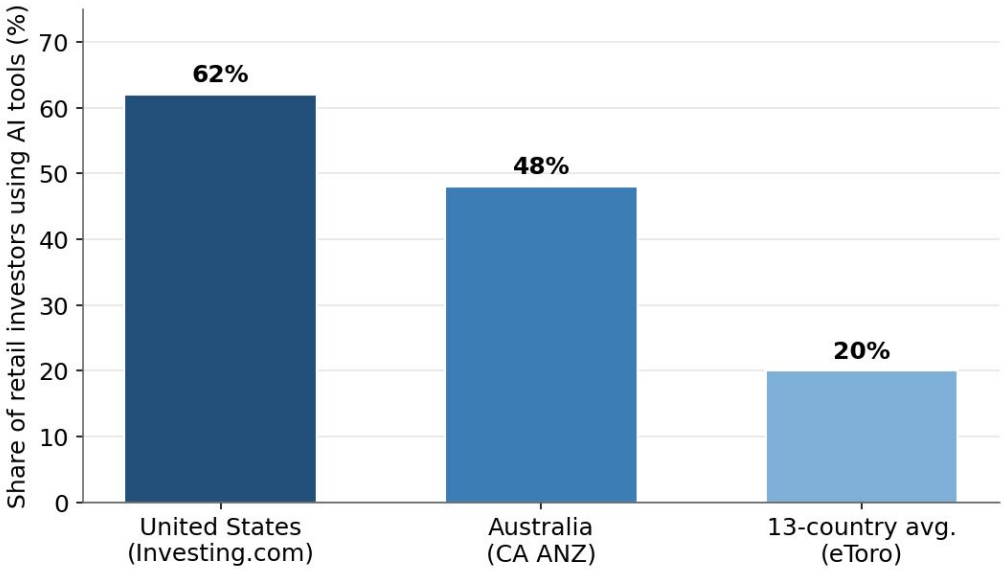
4.1 The Scale and Pace of Adoption

In roughly three years, AI-based investment tools have moved from novelty to default infrastructure. A survey of 938 U.S. retail investors conducted by Investing.com in March 2026 found that 62% used AI to assist with investing decisions — 23.6% regularly, 27.4% occasionally and 11.5% having tried it once or twice — while 21.0% had not yet used AI but were considering it, and 16.6% reported no intention to adopt (Monteiro, 2026). A similar but slightly lower share (48%) of Australian retail investors reported using AI to inform investment decisions in the CA ANZ Australian Investor Confidence Survey 2025, with the highest rate (78%) among those aged 18–29 (van Onselen, 2025). An online survey of about 11,000 retail investors across 13 countries — the eToro 2025 Retail Investor Beat — found that nearly one in five (20%) already use AI tools to make or adjust portfolio decisions (eToro, 2025).

The cross-country pattern shown in Figure 1 reflects both technological-diffusion and language-coverage effects: U.S. and Australian users benefit from more fluent English-language chatbots, while Europe and emerging markets have thinner localised coverage. The pattern is also strongly age-stratified — investors under forty adopt at roughly twice the rate of those over fifty.

4.2 Trust as a Layered Construct

Adoption is not the same as trust. The same Investing.com survey that reported 62% adoption found that only 4% of users “completely” trust AI-generated analysis, 20% “mostly” trust it, 54% trust it “only somewhat” and verify outputs against other sources, and around 23% express little or no trust (Monteiro, 2026). A 2024 Lansons/Opinium survey reported a comparable picture, with only about three in ten U.S. retail investors treating AI-generated

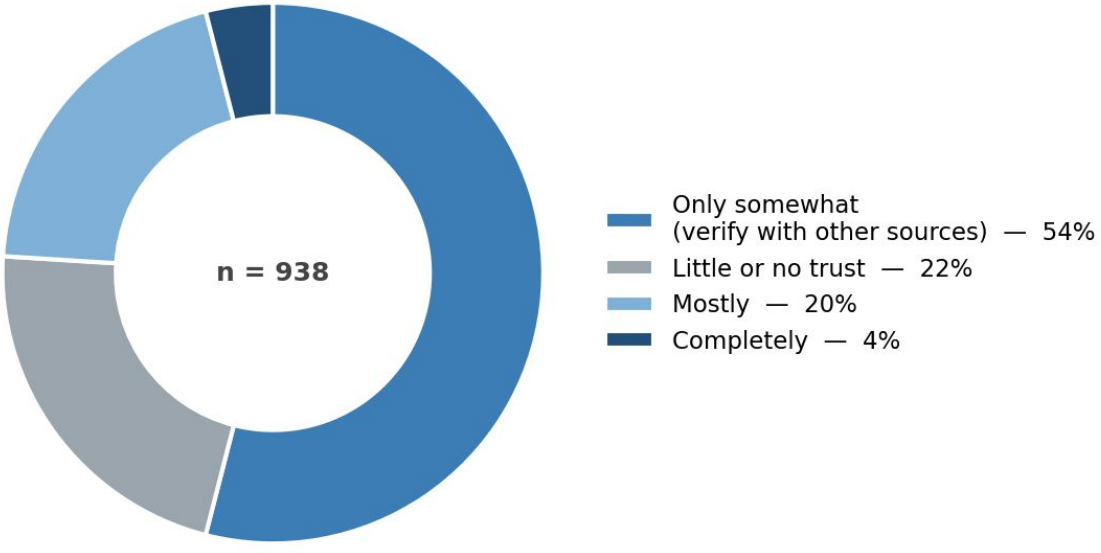


Note: Among Australians aged 18–29, adoption reached 78%. Source: Monteiro (2026); van Onselen (2025); eToro (2025).

Figure 1. Cross-country retail-investor adoption of AI investment tools (2025–2026).

recommendations as a primary input (Lansons, 2024), as shown in Figure 2.

Trust is multidimensional. Among surveyed Australian investors who do not use AI, 43% reported lacking confidence in the underlying information and 46% preferred existing sources — a near-even split between rejecting the tool and being inert about the channel (van Onselen, 2025 This echoes Verma et al. (2025), who found that inertia was the biggest barrier to robo-advisor adoption among Indian investors.



Figures are approximate and sum to ~100%. Source: Monteiro (2026).

Figure 2. Trust in AI-generated investment analysis among U.S. retail investors, 2026.

This pattern indicates that retail investors are using a "trust but verify" heuristic. AI operates more like a speedy first-pass research assistant with outputs checked against outside data than as an oracle. Psychologically, this is different compared to the relationship that most investors currently have with a human advisor, where delegation is deeper but less frequent. To sum up, AI does not so much displace human advice as redistribute cognitive labour: AI is engaged often and trusted lightly; while human advisors are used infrequently but relied upon deeply.

4.3 What Investors Worry About

There are five themes of significant investor concern regarding AI-based recommendations which are ranked in Figure 3 by the share of respondents considering it a major concern. Recommendations that are incorrect or misleading were the most worrying outcome (Phoenix et al. 2019; Zhou et al. 2021), with a level of awareness of the hallucination problem in large language models reflected by being mentioned by an overall average of 38.9% across studies. The second highest priority was thus ranked as fear of a market herding from correlated responses to common signals — the concern that millions of investors who make decisions based on statistically similar AI outputs will trade very strongly correlated (24.2%) and it is particularly heightened among heavy users of AI (Monteiro, 2026). The rest of the concerns involved data privacy, over-dependence that hinders mastering financial-decision making skills, and a short “no worries” category.

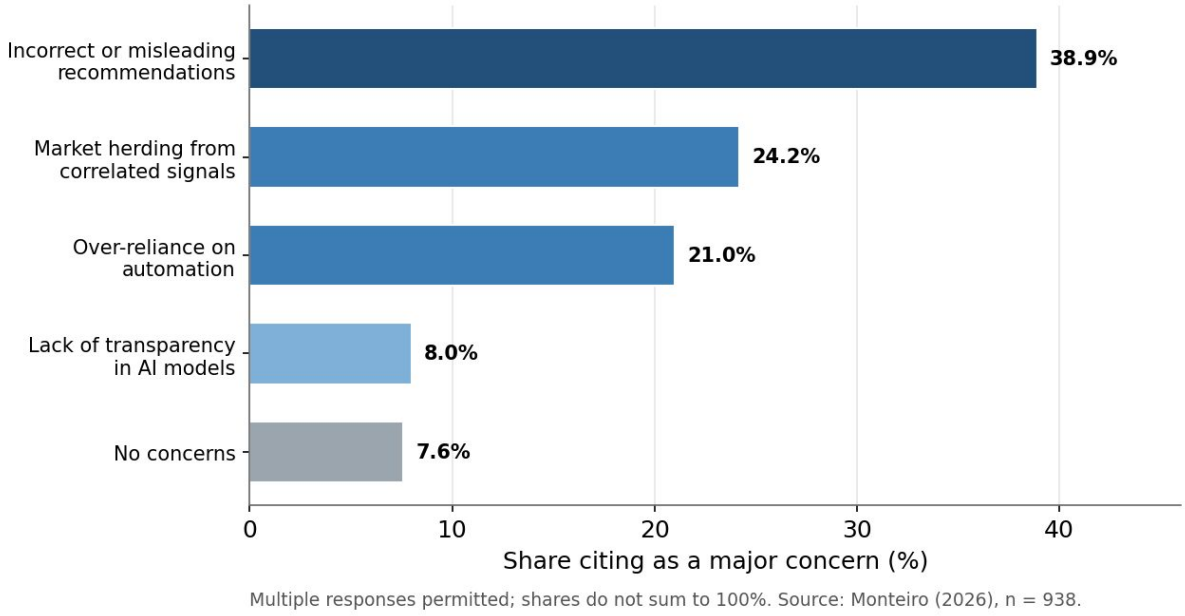


Figure 3. Top concerns about AI-based investment tools among U.S. retail investors.

Hallucination is especially acute for less-followed companies. In a Euronews investigation of ChatGPT-based investment advice, Gaby Diamant, co-founder of the AI investment platform BridgeWise, observed that conversational models generate plausible-sounding but factually unverified content when asked about thinly covered assets (Diamant, 2025). This aligns with Even-Tov, Lourie, Munevar and Nekrasov (2025), who used a natural experiment around Italy’s brief ChatGPT ban and found that retail investors who lost access to generative AI shifted toward popular, simply-valued assets — consistent with the view that AI expands the range of stocks retail investors are willing to hold.

4.4 Effects on Cognitive Biases

The central psychological question is whether AI-based recommendations make retail investors more or less rational. The evidence is mixed but more positive than negative, with important exceptions.

Kulkarni et al. (2025) used structural equation modelling on data from 461 Indian retail investors and found that four perceived qualities of an AI-powered robo-advisor — personalisation, interactivity, autonomy and algorithm interpretability — significantly reduced overconfidence and loss-aversion biases, while the fifth, structural assurance, had no significant effect. A study of NYSE retail investors by

Ahmad et al. (2025) extended this finding, with robo-advisor mediation attenuating biases including availability, anchoring, loss aversion, representativeness and overconfidence. Illustrative path-coefficient magnitudes are shown in Figure 4. Importantly, both studies rely on cross-sectional self-report data and therefore cannot resolve the selection problem raised in Section 2.4: investors who adopt robo-advisors may differ systematically from those who do not, so causal inference about bias reduction is necessarily confounded.

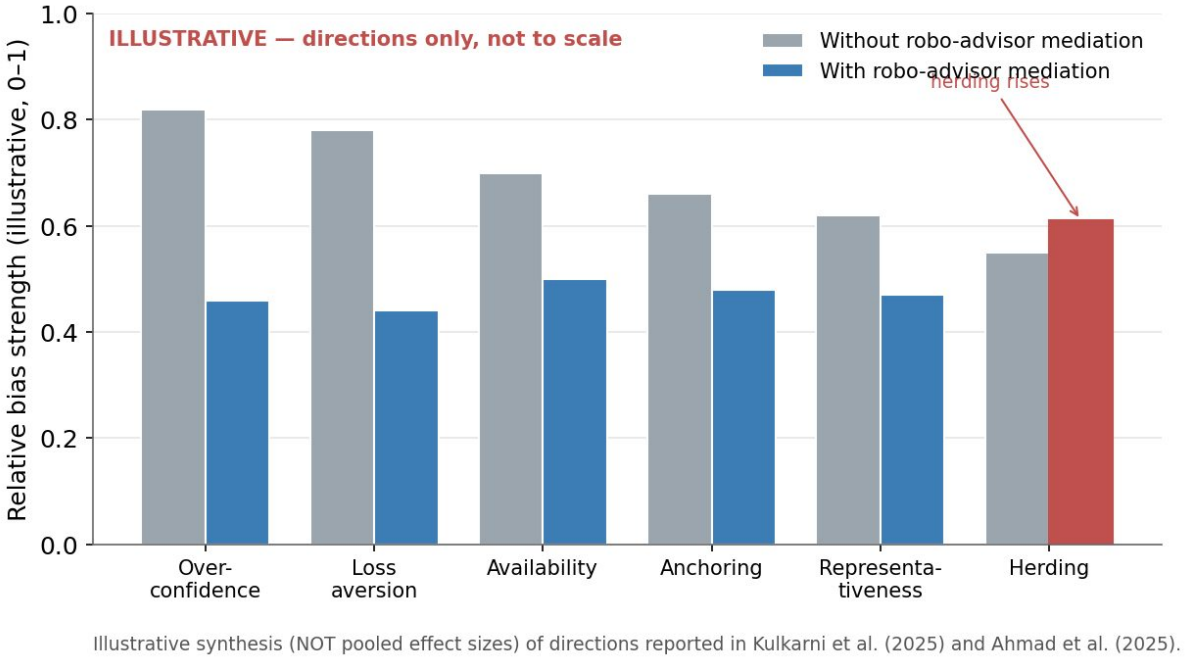


Figure 4. Bias strength among retail investors, with and without robo-advisor mediation (illustrative coefficients synthesised from Kulkarni et al., 2025 and Ahmad et al., 2025; directions only, not pooled effect sizes).

The exception is herding, which robo-advisor mediation appears to strengthen slightly rather than reduce (Figure 4). Hansen and Lee (2025) offer a possible explanation. Running classic herding experiments with large language models, they find that LLM agents acting independently behave more rationally than humans, relying largely on private information rather than market trends — consistent with the bias-mitigation result. But when many investors consult the same or similar AI, the resulting recommendations are highly correlated, producing a new, AI-induced form of herding: commonality between individual psychological processes is replaced by commonality of recommendation. At the level of perception, this concern is echoed in the 2024 Lansons survey, in which a substantial share of U.S. retail investors agreed that AI could produce a herd effect (Lansons, 2024).

4.5 The OSC/BIT Field Experiment

The largest single behavioural-science experiment on this topic to date was conducted by OSC/BIT (2024), which offered 7,771 Canadian residents aged 18 and over a hypothetical CAD 20,000 to invest. Each respondent received an investment recommendation from one of three randomly assigned sources: a human financial-services provider, an AI tool, or a human financial-services provider using an AI tool (the “blended” condition). Adherence was defined as how closely participants followed the allocation they were shown.

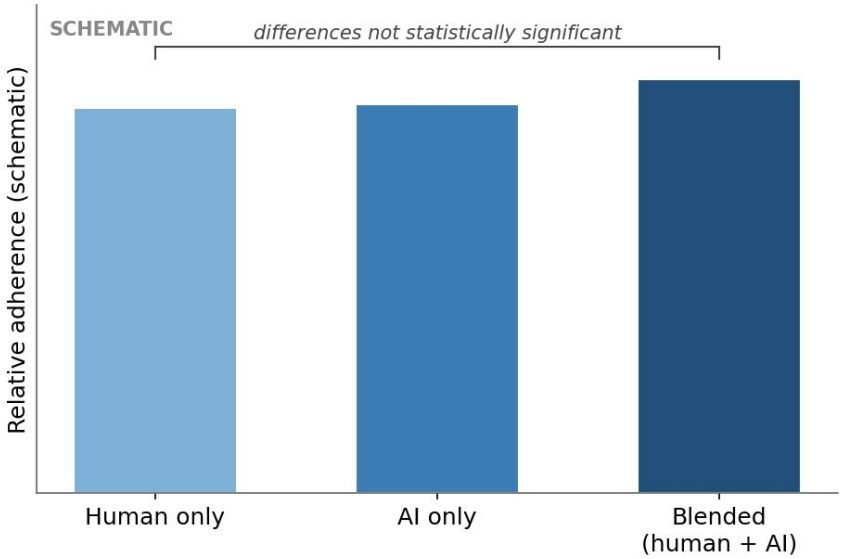


Figure 5. Adherence to investment advice by source of recommendation in the OSC/BIT experiment (n = 7,771). Schematic of the reported pattern; the blended condition adhered most closely, but no condition differed significantly.

Figure 5 shows that adherence was highest in the blended condition and somewhat lower in the human-only and AI-only conditions; however, the pairwise differences were not statistically significant. The absence of significance should not be read as evidence that the conditions are equivalent: the study may simply have been underpowered to detect small effects in subgroup analyses. Even so, the OSC concluded that AI recommendations were no more or less persuasive than human ones and that Canadian retail investors cared little about the source — consistent with treating AI as another channel for distributing financial information. This is more in line with the "algorithm appreciation" of Logg et al. (2019): by 2024, retail investors appear to have largely normalised AI as a valid source of financial information.

Another significant, methodologically interesting finding from the OSC study was that what mattered much more than the source of advice was its quality. And when participants were given low-quality advice — allocations that departed from normative portfolio theory — they adhered to that advice nearly as closely as optimal advice received from either a human or an AI. This does seem to provide some evidence that regulators may be looking at the wrong question in the source-of-recommendation debate which is (human versus AI); rather it may be the quality of the recommendation itself which matters (OSC & BIT, 2024).

4.6 Perceived Performance Impact

Self-reported performance effects are striking. Among the 62% of U.S. retail investors who use AI tools, 65% said AI had improved their market performance — 17% reporting significant improvement and 48% modest improvement. About one-third reported no change, and only 2% reported worse performance (Monteiro, 2026). These figures are shown in Figure 6.

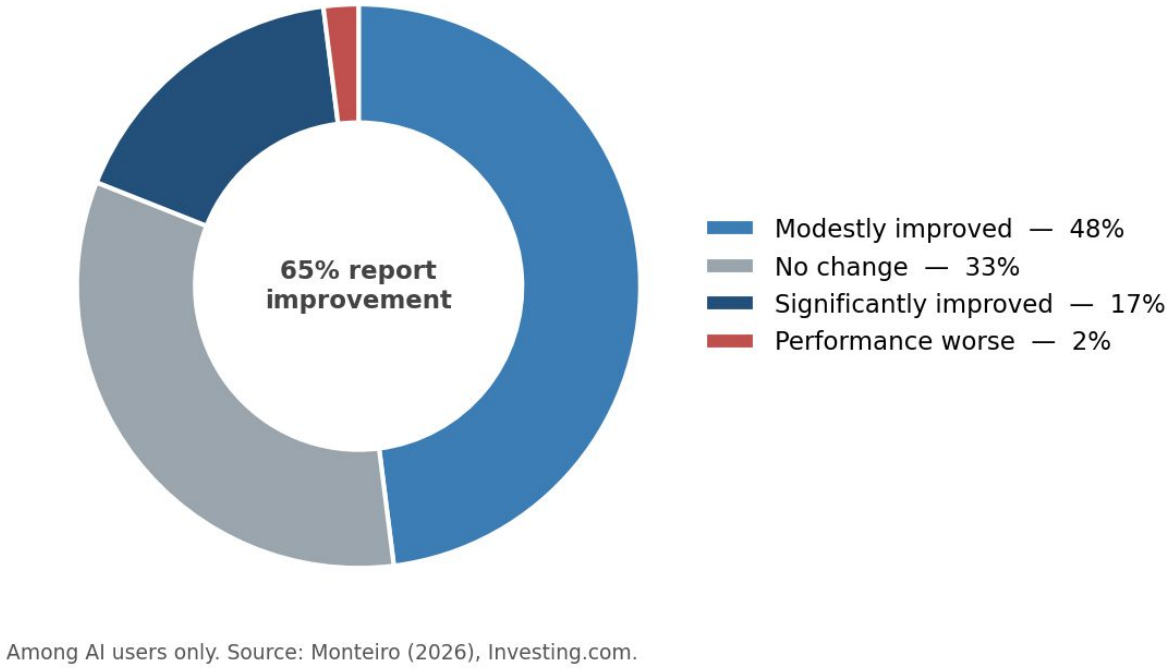


Figure 6. Self-reported impact of AI use on retail-investor performance (Investing.com, 2026).

As always, self-reported performance is subject to confirmation bias and selective memory. Investors are more likely to attribute gains to the tools they chose to use and to discount losses, so the 65% figure should not be treated as an unbiased performance estimate. The account-level brokerage study by Even-Tov et al. (2025) bypasses self-report, even though it addresses a somewhat different question. Using a difference-in-differences design around the Italian ChatGPT ban, the authors show that retail investors concentrated their trading in fewer stocks and were less likely to open new positions — effects that were stronger for illiquid stocks and around earnings announcements. This does not establish a return improvement, nor does it directly support the self-reported performance gains above, but it is at least consistent with the view that AI lowers information-processing costs and expands the investable opportunity set.

4.7 Emerging Psychological Risks

Three psychological risks now appearing in the empirical literature were largely unexplored before 2023. The first is automation bias — the tendency to over-trust automated outputs even when contradicting information is available. In a mixed-methods study published in the Journal of Consumer Behaviour, Bharti, Bharti, Wani and Shahid (2026) found that AI-based product recommendations reduced consumers' sense of accountability for their choices and increased FOMO and dependence on the recommender. Although that study examines consumer products rather than investments, the underlying psychological mechanism is directly relevant: investors who feel they would "miss out" by not following an AI suggestion are more likely to discount evidence that contradicts it.

The second risk is that financial-decision skills erode after years of delegation. Verma et al. The findings of Yang et al. (2025) indicate that positive attitude toward AI reduces resistance to robo-financial by decreasing the necessary of self-reviewing about portfolio. The more pressing issue is that when retail investors outsource scaleable decision making then they are literally outsourcing their own analytical skills and will increasingly become reliant on the very tools whose failures them — who are least equipped to identify flaws in machines over people.

Another threat from the third source is due to output correlation, called herding. Hansen and Lee (2025) Federal Reserve working paper types it pretty clearly in its policy implication: If AI agents make individually rational decisions, then if enough retail investors rely on materially similar models, their trades are more correlated than purely human flows. While this has not

been empirically shown for retail-investor patterns and is rather a theoretical issue, it would therefore be in line with the suspicion that AI trading systems implicitly coordinate in ways to reduce pricing efficiency.

5. Discussion

These findings frame three integrative insights.

The first is that AI is more sensibly regarded as a heuristic rather than oracle. The prevalent survey pattern is one of shallow, repetitive consultation rather than deep delegation of economic calculation to AI. The typical AI investor uses the tool as a speedy first pass research assistant and does not take the results at face value without looking elsewhere for verification. From a psychological perspective, this smashes AI into an availability heuristic with a far longer arm than ever before: it makes ordered and company- and market-specific information feel like instantly accessible information, while the verification stage is theoretically able to preserve some degree of cognitive accountability. This is more optimistic than the automation-bias literature would lead us to expect — but only while investors continue to verify instead of drifting toward passive consumption.

Second, caution called out to bias mediation is evident, yet inconsistent. Evidence from Kulkarni et al. (2025), Ahmad et al. (2025) and Verma et al. (2025) reveals that a properly designed robo-advisor significantly reduces the procyclicality of investor behavior resulting from overconfidence and loss aversion. Finally, two caveats from Section 4.4 should be observed: First, the magnitudes of path coefficients in Figure 4 are correlational, not causal (meaning that we cannot determine their mediation), and second, cross-sectional data cannot test the selection problem — whether robo-advisors harnessed investors who were initially less biased. For herding, the mediation is weakest and the literature on financial-stability considers this as not coincidental; for AI reduces human psychological imitation but creates a new form of technology-imposed correlated behaviour. Since generalizable AI systems can produce dynamics that affect the entire system, regulators and designers of platforms deploying generalizable, human-trained AIs should be conservative in evaluating effects on individual-investor outcomes independent from system-level impacts.

Third thing is that origin doesn't matter that much. This paper from the OSC/BIT (2024) combining the positive and negative influences of peer or social influence among human, AI and blended sources — MAXIM displays broadly similar adherence, but in a finding which has immediate policy implications, that low-quality advice diffuses almost as readily as high-quality advice. The behavioural-protection question is less about "should this advice come from a human or an AI?" than “is this advice good?” This implies moving toward substantive standards for the quality of advice and the data that feed an AI model, as opposed to disclosure-based approaches that merely disclose where a recommendation emanated from.

The integrative framework in Section 3 aids identification of intervention trajectories that are likely to be the most effective. Kulkarni et al. Research by Barocas et al. (2025) shows that interventions at the input layer — changing aspects of the training-data quality, bias causality, and deterministic properties of algorithms for interpretability — have the highest impact on mitigating bias. Mediating-layer interventions to design calibrated trust, wherein investors can adapt their trust in AI commensurate with the actual reliability of those AIs — and outcome-layer interventions that would introduce additional friction into how quickly an investor acts on automated trade ideas generated by the AI — would more precisely hit automation-bias and herding risks.

Central to the normative literature is a tension. The same characteristics that allow AI recommendation systems to be personal, interactive and authoritative are exactly the ones that raise concerns about over-reliance acting as a force for behavioural homogenisation. Investor-protection regulation cannot implicitly avoid this tension and will have to face it squarely.

6. Conclusion

Since the public rollout of large-language-model chatbots, AI-powered investment recommendations have become a commonplace service for retail investors. This paper has compiled the available empirical evidence on how that shift affects the psychology of non-professional investors. Adoption is wide, ranging from roughly one-fifth of investors in some European markets to almost two-thirds in the United States. Trust is layered: most users trust AI only partially and continue to check its outputs, a stance that is less dysfunctional than the automation-bias literature might predict but that depends on continued user vigilance. Well-designed robo-advisors meaningfully

reduce loss-aversion and overconfidence among several classical biases — though this conclusion rests on cross-sectional rather than longitudinal evidence that does not adequately control for self-selection — and they are least effective against herding, where the homogeneity of widely shared AI outputs may itself constitute a new form of herd behaviour.

Using the OSC/BIT 7,771-participant study, we find that retail investors cared little about whether a suggestion was generated by a human, an AI, or a blend, but that they followed low-quality advice almost as closely as high-quality advice. For regulation, this de-emphasises the source of a recommendation in favour of its content quality and the system through which it is delivered.

This synthesis has limitations. Most of the literature it draws on relies on self-report data and spans too few years to capture full market-cycle behaviour, and the evidence base remains geographically concentrated. Because constructs are operationalised differently across studies, an integrative rather than meta-analytic approach was required. Future research should pursue four paths: longitudinal panel studies of retail investors who use AI tools across at least one full market cycle; account-level natural experiments of the kind exploited by Even-Tov et al. (2025); cross-cultural replication of the structural-equation models, given the current concentration of evidence in Indian, U.S., Canadian and Australian samples; and direct empirical testing of the correlated-output herding hypothesis from the Federal Reserve working paper. These questions will only become more urgent as AI systems move from passive recommendation toward agentic execution.

References

Ahmad, U., van Keulen, M., Briassouli, A., & Saad, M. (2025). Cognitive biases, robo advisor and investment decision psychology: An investor's perspective from New York Stock Exchange. Acta Psychologica, 256, 105048. https://doi.org/10.1016/j.actpsy.2025.105048

Barber, B. M., & Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment. Quarterly Journal of Economics, 116(1), 261–292. https://doi.org/10.1162/003355301556400

Bharti, M., Bharti, M., Wani, T. A., & Shahid, S. (2026). The FOMO effect: How AI recommendations drive consumer dependence and weaken brand–self connections. Journal of Consumer Behaviour, 25(3), 1078–1106. https://doi.org/10.1002/cb.70117

Bhatia, A., Chandani, A., & Chhateja, J. (2020). Robo-advisory services in India: A study to analyze the impact of robo-advisor on individuals' personal financial planning. Investment Management and Financial Innovations, 17(4), 75–86.

Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion. Journal of Marketing Research, 56(5), 809–825. https://doi.org/10.1177/0022243719851788

D'Acunto, F., & Rossi, A. G. (2023). Robo-advice: Transforming households into rational economic agents. Annual Review of Financial Economics, 15, 543–563.

Diamant, G. (2025, October). Comments on AI investment advice and hallucination [Interview]. In Personal finance and AI: Should you trust ChatGPT's investment advice? Euronews. [Verify exact title/date.]

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. Journal of Experimental Psychology: General, 144(1), 114–126. https://doi.org/10.1037/xge0000033

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. Management Science, 64(3), 1155–1170. https://doi.org/10.1287/mnsc.2016.2643

eToro. (2025). Retail Investor Beat 2025: Global survey of 11,000 retail investors. eToro Research. [Verify.]

European Securities and Markets Authority. (2025). Public statement on the use of artificial intelligence in the provision of retail investment services. ESMA. [Verify.]

Even-Tov, O., Lourie, B., Munevar, K., & Nekrasov, A. (2025). The effect of AI on retail investor behavior: Evidence from account-level trading data (SSRN Working Paper No. 5778246). https://doi.org/10.2139/ssrn.5778246

Hansen, A. L., & Lee, S. J. (2025). Financial stability implications of generative AI: Taming the animal spirits (Finance and Economics Discussion Series No. 2025-090). Board of Governors of the Federal Reserve System. https://doi.org/10.17016/FEDS.2025.090

Hodge, F. D., Mendoza, K. I., & Sinha, R. K. (2021). The effect of humanizing robo-advisors on investor judgments. Contemporary Accounting Research, 38(1), 770–792. [Verify.]

Jussupow, E., Benbasat, I., & Heinzl, A. (2024). An integrative perspective on algorithm aversion and appreciation in decision-making. MIS Quarterly, 48(4), 1575–1606. [Verify.]

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47(2),

263–291. https://doi.org/10.2307/1914185

Kulkarni, M. S., Patil, K. P., & Pramod, D. (2025). The role of robo-advisors in behavioural finance, shaping investment decisions. Cogent Economics & Finance, 13(1), 2571403. https://doi.org/10.1080/23322039.2025.2571403

Lansons. (2024). U.S. retail investor AI sentiment survey. Lansons in partnership with Opinium. [Verify.]

Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. Organizational Behavior and Human Decision Processes, 151, 90–103. https://doi.org/10.1016/j.obhdp.2018.12.005

Loos, B., Previtero, A., & Scheurle, S. (2020). Robo-advisers and investor behavior (SSRN Working Paper No. 3697469). [Verify.]

Monteiro, T. (2026, April). How retail investors are using AI in 2026 [Survey of 938 U.S. retail investors, conducted March 3–10, 2026]. Investing.com.

Ontario Securities Commission & Behavioural Insights Team. (2024). Artificial intelligence and retail investing: Use cases and experimental research. OSC. https://www.osc.ca/en/investors/investor-research-and-reports/artificial-intelligence-and-retail-investing

Shiller, R. J. (2000). Irrational exuberance. Princeton University Press.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. Science, 185(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

van Onselen, A. (2025). Investor Confidence Survey 2025: Retail investors turn to AI for investment advice. Chartered Accountants Australia and New Zealand (CA ANZ). [Verify.]

Verma, B., Schulze, M., Goswami, D., & Upreti, K. (2025). Artificial intelligence attitudes and resistance to use robo-advisors: Exploring investor reluctance toward cognitive financial systems. Frontiers in Artificial Intelligence, 8, 1623534. https://doi.org/10.3389/frai.2025.1623534