

From Debiasing to Behavioral Reconfiguration: Artificial Intelligence and Retail Investor Decision-Making

Baron William De Rothschild

University of the West of England (UWE Bristol)

Corresponding author: Baron.DeRothschild@uwe.ac.uk

Preprint date: September 2026 | Version: 2.0 (reconciled preprint)

Preprint manuscript - not peer reviewed. Revised SSRN version.

SSRN DOI: 10.2139/ssrn.7534518

Abstract

Artificial intelligence (AI) increasingly shapes retail investment decisions, yet trust, adoption and stated intention are not evidence of changed investment behaviour. This systematic review synthesises a frozen corpus of 99 peer-reviewed studies and distinguishes measured investment outcomes from contextual evidence on acceptance, perceptions and reliance. A documented author-led reassessment of five previously disputed records (M020, M023, M031, M046 and M120) moved them from historical reconfiguration coding to contextual evidence. The resulting descriptive coding distribution is 64 reconfiguration, 11 attenuation, four amplification, three null and 17 contextual records. These are extraction classifications, not verified causal-effect frequencies; the remaining directional records require source-level validation, and no independent second coder was available for the five-case reassessment. The review develops a measurement-sensitive framework connecting system and investor characteristics to cognition, trust, reliance and measurable investment outcomes, while distinguishing conventional robo-advice from interactive generative AI. The framework generates testable propositions rather than claiming that behavioural reconfiguration is empirically dominant.

AI-use disclosure: OpenAI ChatGPT was used to assist with manuscript organisation, consistency checking, language refinement and preparation of a figure. The author is responsible

for the study selection, interpretation, conclusions and final verification. See the full declaration near the end of the manuscript.

Keywords: artificial intelligence; behavioural finance; retail investors; robo-advisors; generative AI; behavioural biases; calibrated reliance

1. Introduction

Artificial intelligence is becoming increasingly embedded in the investment decisions of retail investors. The earlier wave of digital investment advice was dominated by robo-advisory systems that translated investor profiles into automated portfolio recommendations, asset allocations and rebalancing decisions. More recently, conversational and generative AI systems have expanded this role. Investors can now use AI to receive recommendations, interpret disclosures, screen investment opportunities, compare alternatives, ask follow-up questions and refine judgments interactively. This development changes the behavioural-finance problem. In some settings, AI has become part of the process through which investors search for, understand and evaluate financial information, rather than remaining solely an external source of advice.

This change matters because retail investment decisions are already shaped by well-established behavioural mechanisms. Overconfidence, loss aversion, the disposition effect, anchoring, herding and other judgment biases can affect trading, diversification, risk-taking and the interpretation of information. Automated advice has often been presented as a potential corrective because algorithms can apply consistent rules, diversify portfolios and reduce emotion-driven discretion. Evidence from robo-advisory settings supports some of these benefits. D'Acunto, Prabhala and Rossi (2019), for example, document improvements in diversification and reductions in several behavioural distortions after adoption of robo-advice, while Rossi and Utkus (2024) report broader diversification and welfare improvements among users of automated advice.

The emerging literature, however, does not support a simple conclusion that AI makes investors more rational. Some studies find that automated advice reduces particular biases, whereas others report limited or null effects. A growing body of evidence also suggests that AI can generate new

behavioural risks. Investors may avoid useful algorithmic advice, rely on it too readily, respond to anthropomorphic or social cues, or interpret AI-assisted performance as evidence of their own competence. Interactive GenAI makes these tensions especially salient. Croom (2026) shows that greater GenAI interactivity can increase investors' self-assessed ability and willingness to invest without a corresponding improvement in actual information processing. The relevant behavioural question therefore concerns how AI changes the configuration of behavioural mechanisms within the investment decision process, including whether particular biases are reduced.

The nearest reviews require a more precise contribution. Meng, Mahapatra and Xiao (2025) systematically synthesise 89 studies of AI and consumer financial behaviour, including acceptance, interaction, financial behaviour, AI methods and ethical or regulatory concerns, using a TCCM framework. Nourallah et al. (2026) review 50 robo-adviser studies in finance journals and explicitly distinguish asset-management and behavioural-finance streams, including adoption and biased advice; their included papers also touch on ChatGPT-based advice. Accordingly, neither the intersection of AI and financial behaviour, nor robo-advisory bias and adoption, nor the mere inclusion of generative AI can establish novelty here. The intended incremental contribution is a retail-investor-centred, measurement-sensitive synthesis that distinguishes observed choices and portfolio outcomes from trust, attitudes and stated intentions, while examining whether interactive AI changes information processing and reliance. This contribution remains conditional on the completed source audit.

The literature therefore poses three narrower analytical problems. First, empirical studies frequently measure intention, perceived usefulness or trust rather than actual trading, allocation or demonstrated changes in behavioural bias; these evidence types must not be treated as interchangeable. Second, studies with measured investment outcomes need a consistent distinction among outcome, comparator, investor population and study design before directional patterns can be inferred. Third, the emerging distinction between conventional automated recommendations and interactive generative AI warrants a mechanism-based comparison, but only where primary studies directly support the proposed pathways.

Despite this progress, three limitations remain. First, much of the literature is organised around adoption, intention to use or broad AI/FinTech capability rather than the behavioural consequences of AI use during retail investment decisions. This leaves post-adoption processes

such as reliance, delegation, information processing and portfolio execution comparatively fragmented. Second, behavioural effects are usually examined separately: a technology may reduce one bias while leaving another unchanged or amplifying a different distortion, yet these outcomes are rarely integrated within a common behavioural framework. Third, conventional robo-advice and interactive GenAI are often grouped within broad AI-enabled finance even though their behavioural roles differ. Traditional robo-advice generally produces an external recommendation; interactive GenAI can enter the investor's cognitive process through repeated dialogue, simplification and adaptive information search.

This review addresses these limitations by synthesising the behavioural evidence on AI-enabled and algorithmic investment technologies from the perspective of retail investor decision-making. Rather than asking only whether investors adopt AI, the review examines how AI affects behavioural biases, trust, risk perceptions, information processing, reliance, delegation and investment outcomes. It also distinguishes the traditional automated-advice model from the emerging interactive GenAI model and examines the conditions under which AI attenuates, amplifies, reconfigures or leaves unchanged behavioural influences on investment decisions.

The review is guided by four research questions:

RQ1. How are AI-enabled technologies used in retail investment decision-making, and how has this evolved from automated robo-advice to interactive and generative AI?

RQ2. Which behavioural biases and decision-making mechanisms are associated with AI-assisted retail investing, and what evidence exists that AI attenuates, amplifies or reconfigures them?

RQ3. How do investor characteristics, AI characteristics and decision contexts influence these behavioural effects?

RQ4. What integrative behavioural-finance framework can explain the relationship between AI-assisted decision-making, behavioural mechanisms and investor outcomes?

The review makes four conditional contributions: it maps the frozen 99-study corpus while separating adoption and trust evidence from measured behavioural outcomes; documents an auditable five-case author-led coding correction and sensitivity analysis; compares conventional robo-advice with interactive generative AI at the level of investor decision processes; and

proposes an integrative framework and testable research agenda. The original coding distribution is preserved as an archival baseline, not presented as verified effect prevalence. Independent second-reviewer adjudication was unavailable, and other directional classifications retain source-verification limitations.

These contributions are developed through an integrative conceptual framework linking AI/platform and investor/context characteristics to cognitive and affective mechanisms, reliance and interaction states, behavioural consequences and investment outcomes. The framework is complemented by a TCCM-based future research agenda that prioritises GenAI cognitive partnership, calibrated trust, multi-bias behavioural reconfiguration, human-AI design and investor welfare, investor heterogeneity and vulnerability, and the connection between AI advice, execution and realised financial outcomes.

The remainder of the article is organised as follows. Section 2 develops the behavioural-finance and human-algorithm foundations of the review. Section 3 describes the systematic review methodology. Section 4 presents the descriptive profile of the evidence base. Section 5 develops the six-theme thematic synthesis, followed by the integrative conceptual framework in Section 6. Section 7 discusses theoretical, investor, platform and regulatory implications. Section 8 presents the future research agenda. The final sections address the limitations of the review and conclude with the main implications for behavioural finance and AI-assisted retail investing.

2. Theoretical Background

This review is positioned within behavioural finance rather than within technology adoption alone. The central question concerns how automated and interactive systems change the behavioural processes through which retail investment decisions are formed after investors begin using AI-enabled tools. This distinction is important because the same technology can alter information processing, perceived competence, trust, delegation and emotional responses before any observable portfolio decision is made. The theoretical background therefore begins with established behavioural-finance mechanisms, then considers how algorithmic and AI-enabled advice changes the decision environment in which those mechanisms operate.

2.1 Behavioural finance and retail investor decision-making

Behavioural finance developed in response to evidence that investors do not always make decisions in the manner implied by fully rational models. Prospect theory established that decisions under risk depend on reference points and that losses can carry greater psychological weight than equivalent gains (Kahneman and Tversky, 1979). In investment settings, such asymmetry is reflected in behaviours such as loss aversion and the disposition effect. Shefrin and Statman (1985) conceptualised the tendency to sell winning investments too early while continuing to hold losing positions, and subsequent transaction-level evidence documented this behaviour among individual investors (Odean, 1998).

The relevance of these mechanisms extends beyond the disposition effect. Retail investors can exhibit overconfidence, anchoring, herding, confirmation tendencies and biased responses to recent or salient information. These behaviours affect trading frequency, risk-taking, diversification, security selection and the interpretation of financial information. Behavioural biases need not operate independently. An investor can simultaneously be loss averse, overconfident in personal judgment and influenced by social information. This makes the behavioural consequences of AI theoretically ambiguous: a system may suppress one source of bias while leaving another unchanged or creating a new route through which confidence or reliance becomes distorted.

For this reason, the present review treats behavioural bias as a dynamic feature of the decision process rather than as a fixed characteristic of the investor. The relevant question is whether AI changes the strength, expression or location of behavioural influence. A recommendation system may reduce emotional discretion at the point of portfolio rebalancing, while an interactive system may increase confidence earlier in the information-processing stage. The behavioural effect can therefore depend on where AI enters the investment process.

2.2 Automated advice as a behavioural intervention

Robo-advisory systems introduced a new form of investor decision support by converting information about investor objectives, risk preferences and constraints into automated portfolio recommendations. In behavioural terms, this process can operate as a commitment or structuring device. By formalising risk profiling, diversification and rebalancing, automated advice can reduce opportunities for emotion-driven or heuristic trading. Evidence from the reviewed

literature supports this possibility. D'Acunto, Prabhala and Rossi (2019) found that adoption of robo-advice improved diversification among initially under-diversified investors and was associated with declines in the disposition effect, trend chasing and rank-related trading behaviour.

However, the behavioural value of automation depends on whether investors actually use the recommendation and how much decision authority they are prepared to delegate. Automated advice therefore introduces a second behavioural problem: the investor must decide whether to rely on an algorithm. This creates the possibility of both under-reliance and over-reliance. A technically sound recommendation can produce little benefit if an investor rejects it because it is algorithmic, while a weak or inappropriate recommendation can become harmful if it is followed uncritically.

This dual possibility means that robo-advice should not be conceptualised only as a mechanism for replacing biased human judgment with objective computation. The technology also becomes a new object of judgment. Investors evaluate its competence, transparency, autonomy, institutional backing and similarity to human advice. Behavioural outcomes therefore depend on the quality of the algorithm and on how investors perceive and interact with it.

2.3 Trust, algorithm aversion and algorithm appreciation

Research on human reliance on algorithms demonstrates that reactions to algorithmic advice are not uniformly sceptical or favourable. Dietvorst, Simmons and Massey (2015) showed that people can become disproportionately reluctant to use an algorithm after observing it make an error, even when the algorithm performs better than a human forecaster. This phenomenon, described as algorithm aversion, is relevant to financial advice because investment recommendations are inherently uncertain and visible errors may rapidly reduce confidence in automated systems.

At the same time, algorithm aversion is not universal. Logg, Minson and Moore (2019) found evidence of algorithm appreciation, with participants placing greater weight on advice when they believed it came from an algorithm rather than a person in several judgment tasks. Taken together, these two bodies of evidence suggest that investor reliance should not be represented as a simple preference for either human or machine advice. Reliance can change with experience,

observed errors, perceived expertise, decision stakes and the degree of control retained by the user.

Trust is therefore theoretically important because it links perceptions of the technology to behavioural reliance. Yet the distinction between trust and appropriate reliance is essential. High trust may encourage useful delegation when the system is accurate, but it may also reduce scrutiny. Low trust may protect an investor from poor advice, but it can also prevent use of beneficial automation. The concept of calibrated reliance used later in this review reflects this tension. It refers to reliance that is proportionate to system quality, uncertainty and task demands, rather than to the maximisation of trust itself.

2.4 Human-AI interaction and the social presentation of advice

Automated advice is also shaped by the way it is presented. Financial decision-making is not purely informational; source cues, social presence and perceptions of competence can influence whether advice is accepted. Conversational interfaces, anthropomorphic features and human-in-the-loop arrangements can make an automated system appear more socially responsive or credible. Within the reviewed literature, conversational robo-advisors have been associated with higher affective trust and greater recommendation acceptance, while anthropomorphic or competence-related cues can alter perceived trustworthiness and willingness to follow investment advice.

These design effects create a further theoretical complication. Humanisation can increase engagement without necessarily improving decision quality. If social cues encourage investors to trust advice for peripheral reasons, the interface may substitute one behavioural heuristic for another. Conversely, human support may help investors interpret uncertainty and maintain appropriate control in complex or high-stakes decisions. Human-AI interaction should therefore be treated as a behavioural mechanism rather than as a cosmetic feature of platform design.

2.5 From robo-advice to interactive generative AI

Generative AI represents a more substantial change in the investor decision environment than conventional robo-advice. Traditional robo-advisory systems generally follow a relatively structured sequence: the investor provides information, the system processes that information, and a recommendation or portfolio is returned. The central behavioural questions concern trust,

acceptance, delegation and execution. Interactive GenAI changes this sequence because the investor can engage in repeated dialogue with the technology, ask follow-up questions, request explanations, compare alternatives and refine the information being considered.

This interactivity means that GenAI can become part of the investor's reasoning process rather than remaining an external source of advice. It may reduce cognitive load, simplify disclosures, contextualise financial information or help investors formulate questions. At the same time, it can blur the boundary between the capabilities of the user and those of the system. Croom (2026) provides direct evidence of this risk: across three experiments, greater GenAI interactivity increased investors' self-assessed ability and willingness to invest even though it did not improve, and in some conditions hindered, actual information processing.

The distinction is theoretically significant because it introduces metacognition into the analysis of AI-assisted investing. With traditional robo-advice, an investor may ask whether the algorithm is competent. With interactive GenAI, the investor may also update beliefs about personal competence because the system has made a difficult task feel easier. This creates the possibility of cognitive assistance and metacognitive misattribution occurring at the same time. Accordingly, the behavioural effect of GenAI cannot be assessed only by measuring recommendation quality or adoption. It also requires attention to information processing, confidence, cognitive offloading, independent verification and repeated interaction.

2.6 Theoretical positioning of the review

The preceding literature provides three foundations for the review. First, behavioural-finance theory establishes that investor decisions are shaped by systematic biases and affective responses rather than by information alone. Second, research on algorithmic judgment shows that introducing an algorithm creates new questions of trust, aversion, appreciation and delegation. Third, interactive GenAI extends the decision environment by becoming part of the process through which investors search for, interpret and evaluate information.

These foundations imply that existing technology-adoption models, while useful for explaining intention to use, are insufficient on their own for the present research problem. Constructs such as perceived usefulness, ease of use and social influence can explain whether an investor adopts a system, but they do not fully explain whether AI subsequently reduces loss aversion, increases overconfidence, changes reliance, alters information processing or affects portfolio execution.

The review therefore uses adoption-oriented theories as supporting perspectives while keeping behavioural-finance mechanisms and investment outcomes at the centre of the analysis.

On this basis, the review does not begin with the assumption that AI is either rationalising or behaviourally harmful. Instead, it examines whether AI attenuates existing biases, amplifies behavioural distortions, reconfigures the relationship among biases and decision mechanisms, or produces no material behavioural effect. This theoretical position provides the basis for the systematic review methodology in Section 3 and for the thematic synthesis and conceptual framework developed later in the article.

3. Review Methodology

3.1 Review design

Evidence-audit status: the 99-study inclusion set is frozen. The author reassessed M020, M023, M031, M046 and M120 against previously documented full-text outcome evidence and classified all five as contextual adoption, trust or intention studies rather than demonstrated investment-behaviour reconfiguration. The original codes are preserved in the archive; the revised descriptive distribution is reported with a five-case sensitivity comparison. This was author-led reassessment, not independent adjudication. Other directional classifications have not all undergone complete primary-source endpoint verification, so the counts must not be interpreted as validated causal-effect prevalence. Journal-specific requirements and publication declarations remain separate submission tasks.

This study adopted a systematic literature review approach to examine how artificial intelligence-enabled investment technologies influence behavioural mechanisms and investment decision-making among retail investors. The review was designed to move beyond a conventional technology-adoption perspective by focusing specifically on the behavioural processes through which automated and AI-enabled investment technologies may attenuate, amplify or reconfigure investor biases and decision-making.

The review process was guided by the principles of the PRISMA 2020 framework to support transparency in study identification, screening, retrieval and inclusion (Page et al., 2021). The protocol was established before full-text synthesis and specified the review period, databases, eligibility criteria, screening procedure, extraction fields and synthesis strategy.

The review focused on peer-reviewed research published between 2015 and 2026. The starting point of 2015 was selected to capture the period in which algorithmic and robo-advisory technologies became increasingly relevant to retail investment decision-making, while also accommodating the more recent transition towards conversational and generative AI. Earlier foundational behavioural-finance literature was retained for theoretical background where necessary but was not included within the contemporary systematic-review sample unless it directly examined technology-assisted investment decision-making.

The review was structured around four research questions:

RQ1. How are AI-enabled technologies used in retail investment decision-making, and how has this evolved from automated robo-advice to interactive and generative AI?

RQ2. Which behavioural biases and decision-making mechanisms are associated with AI-assisted retail investing, and what evidence exists that AI attenuates, amplifies or reconfigures them?

RQ3. How do investor characteristics, AI characteristics and decision contexts influence these behavioural effects?

RQ4. What integrative behavioural-finance framework can explain the relationship between AI-assisted decision-making, behavioural mechanisms and investor outcomes?

3.2 Eligibility criteria

Eligibility criteria were defined to maintain a clear behavioural-finance and retail-investor focus. Studies were eligible when they examined individual, retail, household, private or non-professional investors, or participants making individual investment decisions. The technology examined had to form a substantive part of the investor decision environment and could include robo-advisors, automated or algorithmic investment advice, AI-enabled investment platforms, conversational AI, generative AI, large language models, or related forms of machine-supported financial advice.

For inclusion, studies also needed to contain a meaningful behavioural or investment-decision component. Relevant constructs included behavioural biases, trust, algorithmic reliance, algorithm aversion or appreciation, automation bias, perceived risk, financial literacy, decision

confidence, delegation, information processing, cognitive offloading, investment judgment, portfolio choice or related behavioural outcomes.

Both empirical and theoretically relevant conceptual journal articles were initially eligible for screening. The core synthesis, however, prioritised studies providing direct investor-level evidence concerning investor-facing AI or algorithmic advice. Empirical methods could include experiments, surveys, archival or administrative data, field studies, qualitative research and mixed-method designs.

Studies were excluded when their primary focus concerned institutional quantitative trading, high-frequency trading, market making, purely technical stock-price prediction, fraud detection, credit scoring, anti-money-laundering applications or other uses of AI where the technology did not directly enter an individual investor's decision environment. Studies employing machine learning only as a researcher-side analytical technique were similarly excluded from the behavioural core. Research centred on firms, professional asset managers or institutional decision-makers was excluded unless direct evidence concerning individual investment behaviour was also present.

Only peer-reviewed journal articles published in English between 2015 and 2026 were included in the contemporary evidence base. Conference papers, dissertations, working papers, blogs, professional reports and other non-peer-reviewed materials were not included in the core synthesis.

3.3 Search strategy and information sources

The formal database search was conducted using Scopus and Business Source Ultimate. These databases were selected to provide coverage across finance, behavioural research, business, information systems and financial-technology scholarship.

The search strategy was developed around three conceptual blocks: (1) AI-enabled investment technology, including terms relating to robo-advisors, algorithmic advice, automated financial advice, generative AI and conversational AI; (2) the investor and investment-decision context, including retail investors, individual investors, investment decisions, portfolio choice and investment advice; and (3) behavioural mechanisms, including behavioural bias, overconfidence,

disposition effect, herding, loss aversion, trust, risk perception, algorithm aversion, automation bias, delegation and financial literacy.

The Scopus search was conducted within titles, abstracts and keywords. The final search incorporated separate terminology for conventional robo-advisory technologies and newer forms of generative and conversational AI. Publication-year restrictions were applied to retain studies published from 2015 to 2026. Results were subsequently restricted to journal articles and reviews published in English.

The initial Scopus search produced 366 records. Following application of the document-type, language and source filters, 202 records remained. The corresponding Business Source Ultimate search returned 117 records. The full-text availability filter was not applied because doing so could have systematically excluded relevant studies based on institutional access rather than scholarly eligibility. Immediately before submission, the database searches were updated on 5 September 2026 using the same eligibility logic. The update returned 18 Scopus records and 8 Business Source Ultimate records. Business Source Ultimate permitted month/year filtering only, so the update window was set to January-September 2026 and executed on 5 September 2026.

Web of Science Core Collection was considered during protocol development but could not be searched through the available institutional access. It is therefore not reported as a searched database.

3.4 Deduplication and screening

Across the original and September-update searches, 345 database records were identified. Deduplication used DOI information where available, followed by normalised-title comparison and manual checking. In total, 61 duplicate records were removed before screening, leaving 284 unique records for title-and-abstract assessment.

A total of 284 unique records entered title-and-abstract screening. Screening was conducted against the predefined inclusion and exclusion criteria, with particular attention to whether AI was situated within the investor's decision environment rather than being used solely as a researcher-side analytical tool. Studies were initially classified as include, exclude or uncertain. Records retained following title-and-abstract screening proceeded to full-text retrieval.

Of the 284 records screened, 158 reports were sought for full-text retrieval and 126 records were excluded at the title-and-abstract stage.


Eight reports could not be retrieved despite further access attempts. These were recorded as reports not retrieved, and no eligibility decision was inferred from their abstracts. The additional not-retrieved report identified in the September update was Mohapatra et al. (2026), for which institutional retrieval was unsuccessful and an author copy was requested. Consequently, 150 full-text reports were assessed for eligibility.

Following full-text assessment, 51 reports were excluded because they did not satisfy the final eligibility requirements. This resulted in a final core evidence base of 99 studies at the 5 September 2026 review cut-off.

The final study-selection sequence was: 345 records identified; 61 duplicates removed; 284 records screened; 158 reports sought for retrieval; 8 reports not retrieved; 150 full-text reports assessed; 51 full-text exclusions; and 99 studies included. These stages are reported in the updated PRISMA flow diagram in Figure 1.

Figure 1 presents the PRISMA 2020 flow diagram for the identification, screening, retrieval and inclusion stages used in this review.

Figure 1. PRISMA 2020 flow diagram for study identification and selection



Counts combine the original search and the 5 September 2026 update: 345 records identified, 61 duplicates removed, 284 title/abstract records screened, 158 reports sought, 150 assessed, 51 excluded after full-text review, and 99 studies included.

Figure 1. PRISMA 2020 flow diagram for study identification and selection. The visual rendering of this figure was prepared with assistance from OpenAI ChatGPT using the author-defined and verified PRISMA screening counts. The review methodology, screening decisions, numerical counts and interpretation were determined and verified by the author.

3.5 Data extraction

A structured extraction matrix was developed to ensure consistent coding across the included studies. For each study, bibliographic information and substantive methodological and behavioural characteristics were recorded.

The extraction fields included study design and theoretical lens; method and sample; geographical context; AI or investment technology examined; role played by AI within the decision process; investor type; behavioural bias or mechanism; trust-related factors; algorithmic reliance; risk perception; financial literacy; main findings; moderators and mediators; limitations; research gaps; and overall evidence strength. Complete study-level bibliographic and coding details for all 99 included studies are provided in Supplementary Table S1.

The role of AI was additionally classified according to its position within the investor decision process: information provider, decision support, delegated adviser, interactive cognitive partner, automated executor, hybrid human-in-the-loop, or other. This distinction was important because the behavioural implications of a conventional automated portfolio recommendation may differ materially from those arising when an investor engages iteratively with a generative AI system.

3.6 Assessment of evidence strength

The included studies varied considerably in research design and evidential strength. To prevent findings from being treated as equally persuasive, each study was assigned an evidence-strength assessment of High, Moderate or Low.

Assessment considered the appropriateness of the research design, sample characteristics, validity of the measures, analytical rigour and the extent to which the reported conclusions were supported by the underlying methodology.

Experimental studies, field experiments, large-scale administrative datasets and well-identified observational designs generally provided stronger evidence for behavioural effects than small convenience surveys or purely correlational self-report studies. However, methodological form alone did not determine the classification; the appropriateness and execution of the individual study were also considered.

Across the final 99-study sample, 32 studies were assessed as High evidence strength, 61 as Moderate and 6 as Low. The evidence-strength assessment was used interpretively rather than as

a mechanical exclusion threshold. Lower-strength studies were retained where they contributed relevant evidence, but stronger causal claims in the synthesis were weighted towards studies with more robust research designs.

3.7 Coding and thematic synthesis

The synthesis followed a staged coding process. First, descriptive coding captured constructs explicitly examined by each study, including overconfidence, herding, loss aversion, disposition effect, trust, financial literacy, algorithm aversion, automation bias, anthropomorphism, delegation and information processing.

Second, interpretive coding considered what the findings implied about the behavioural role of AI. Particular attention was given to whether AI appeared to reduce an existing behavioural distortion, intensify it, leave it unchanged or alter the form through which it influenced investment behaviour.

For directional synthesis, attenuation, amplification, reconfiguration and null effects were considered only where an investor-level behavioural or investment outcome was supported by the reported endpoint. Adoption, trust, recommendation and future-use intention were coded as contextual mechanisms unless a study separately measured investment choice or behavioural change. Historical extraction labels are preserved for auditability; the five-case author-led reassessment is reported separately from independent adjudication, which was unavailable.

Third, the descriptive and interpretive codes were consolidated into higher-order themes. Initial thematic assignments were subsequently subjected to a manual thematic audit in which boundary cases, overlapping constructs and contradictory findings were reconsidered against the full extracted evidence.

This process produced six final higher-order themes: (1) selective debiasing and portfolio discipline; (2) behavioural amplification and AI-induced reconfiguration; (3) trust, resistance and calibrated reliance; (4) human-AI interaction, social cues and hybrid advice; (5) investor heterogeneity, literacy and vulnerability; and (6) information processing, delegation and GenAI cognitive partnership.

Contradictory findings were retained rather than resolved through simple majority interpretation. For example, evidence showing that robo-advice mitigated particular behavioural biases was

considered alongside studies reporting null effects or amplification of different biases. These contradictions subsequently informed the conceptual distinction between behavioural attenuation, amplification and reconfiguration.

Five-case author reassessment and sensitivity analysis. Previously documented full-text outcome checks for M020, M023, M031, M046 and M120 identified continued-use or recommendation intention, trust formation/transfer, prospective adoption intention, platform-use intention, and adoption attitudes/trust, respectively. None of those endpoints, as extracted, establishes an AI-caused change in actual trading, allocation or investment bias. The author therefore reclassified all five from historical reconfiguration to contextual evidence, retaining them in the 99-study corpus and preserving their original codes in the audit archive. No independent second coder completed adjudication; invited colleagues were unavailable. The original and revised descriptive counts are presented as a sensitivity analysis, not estimates of causal-effect prevalence. Other directional records remain subject to source-level verification.

3.8 Development of the conceptual framework and future research agenda

The conceptual framework was developed inductively from the thematic synthesis rather than being specified before the review. Relationships repeatedly observed across the six themes were organised into a sequential architecture comprising AI/platform and investor/context antecedents, cognitive and affective mechanisms, reliance and interaction states, behavioural consequences, and investment outcomes. Investor characteristics and decision context were retained as boundary conditions rather than assuming homogeneous effects across retail investors.

A specific distinction was also made between conventional robo-advisory systems and interactive generative AI. Traditional robo-advice generally follows a profile-algorithm-recommendation structure, whereas generative AI allows repeated interaction in which investors can use AI as part of the information-processing and judgment process itself.

Future research directions were subsequently organised using a Theory-Context-Characteristics-Methodology (TCCM) logic. Proposed research questions were included only where they could be linked to an observed evidential gap, contradiction or methodological limitation identified during the review.

4. Descriptive Results

4.1 Profile and growth of the evidence base

The final review comprised 99 studies published between 2019 and 2026. Although the eligibility window began in 2015, no study published between 2015 and 2018 met the criteria for inclusion in the final core sample. The evidence base expanded sharply in the later years of the review period. Only seven included studies were published between 2019 and 2022, compared with nine in 2023, 14 in 2024, 29 in 2025 and 35 in 2026. Consequently, 64 studies (64.6%) were published in 2025 or 2026, and 78 (78.8%) were published from 2024 onward.

The increase in publication volume was accompanied by a broadening of the technologies examined. The sample remained dominated by robo-advisory and automated investment-advice systems, but more recent studies increasingly examined generative AI, conversational interfaces, AI-mediated financial information, predictive tools and other forms of interactive decision support. This temporal shift is important because it suggests a movement from research centred primarily on whether investors adopt automated advice towards research examining how investors interact with AI during information processing, judgment and portfolio decision-making.

4.2 Geographical distribution

The included evidence was geographically diverse, although it was not evenly distributed. India was the most frequently represented single-country setting, with at least 24 studies, followed by the United States with at least 14 and China with 10. Indonesia contributed 4 studies, while South Korea contributed 3. Saudi Arabia, the United Arab Emirates and Pakistan were also represented by multiple studies. A smaller number of studies examined investors in Malaysia, Sweden, Thailand, Turkey, Taiwan, Vietnam, Japan, New Zealand and the Netherlands, alongside cross-country or platform-based samples.

This distribution shows substantial representation from emerging and rapidly digitalising financial markets, particularly India and wider Asia. At the same time, several records did not report a sufficiently precise geographical location in the available study information. Geography was therefore retained as reported rather than inferred. The geographical pattern also reinforces the need to treat institutional setting, market maturity and cultural context as potential boundary conditions rather than assuming that evidence obtained in one market transfers directly to all retail investors.

4.3 Research designs and evidence strength

Cross-sectional survey and structural-model designs represented the largest methodological group in the sample, accounting for 57 studies (57.6%). Experimental designs, including laboratory, online, field, discrete-choice and multi-study experiments, accounted for 23 studies (23.2%). A further 8 studies used archival, transaction, administrative or observational field data, while 6 employed mixed methods and 3 were primarily qualitative. The remaining studies used other empirical combinations that did not fit cleanly within these broad categories.

The evidence-strength assessment reflected this methodological heterogeneity. Of the 99 studies, 32 (32.3%) were assessed as High evidence strength, 61 (61.6%) as Moderate and 6 (6.1%) as Low. The predominance of Moderate-strength evidence is consistent with the large number of cross-sectional survey studies. However, the growing presence of experiments, platform data, field evidence and administrative datasets provides an increasingly strong basis for examining actual behavioural responses rather than adoption intentions alone.

4.4 Technologies and the role of AI

Robo-advisory and automated investment advice remained the dominant technological context, but the role played by AI varied materially across studies.

A more informative distinction concerned the role AI played within the investor decision process. Decision support was the largest category, with 62 studies (62.6%). Interactive cognitive-partner applications accounted for 13 (13.1%), hybrid human-in-the-loop arrangements for 10 (10.1%), information-provider roles for 8 (8.1%), delegated-adviser roles for 4 (4.0%), and automated-executor roles for 2 (2.0%). This pattern shows that the literature increasingly extends beyond a simple automated-advice model and includes settings in which AI interacts with investors, supports interpretation, shares decision authority or influences execution.

4.5 Direction of reported behavioural effects

The frozen extraction originally coded 69 studies as reconfiguration, 11 as attenuation, four as amplification and three as null, with 12 contextual/not direction-coded records (total n=99). Author-led full-text endpoint reassessment of five disputed records (M020, M023, M031, M046 and M120) identified adoption, intention or trust endpoints rather than demonstrated investment-behaviour change. Their revised classification is contextual, yielding 64 reconfiguration, 11

attenuation, four amplification, three null and 17 contextual records (n=99). The revised figures are descriptive coding counts and should not be interpreted as the verified distribution or causal prevalence of AI-induced effects. Independent second-reviewer adjudication was not obtained, and the remaining directional records have not all completed source-level verification.

The five-case sensitivity analysis changes reconfiguration coding from 69 to 64 and contextual coding from 12 to 17 without changing inclusion (n=99). The apparent predominance of reconfiguration in either coding version is not proof that AI commonly reconfigures observed investor behaviour. The synthesis distinguishes (i) executed trading, allocation and directly measured bias outcomes, (ii) incentivised or hypothetical experimental choices, and (iii) trust, attitudes, adoption and stated intentions. The last category informs mechanisms and boundary conditions, not behavioural-effect direction.

Sensitivity note: the original five records were coded as reconfiguration in the archived matrix; their revised contextual status reflects endpoint-based author reassessment, not exclusion or independent dual coding. Detailed case rationales and source identifiers are recorded in the companion five-case adjudication workbook.

Table 1. Descriptive profile of the 99 included studies

Dimension	Category	n	% of 99
Publication year	2025–2026	64	64.6
Publication year	2024–2026	78	78.8
Method	Cross-sectional survey / model	57	57.6
Method	Experimental	23	23.2
Method	Archival / field / observational	8	8.1
Evidence strength	High	32	32.3
Evidence strength	Moderate	61	61.6
Evidence strength	Low	6	6.1
Role of AI	Decision support	62	62.6
Role of AI	Interactive cognitive partner	13	13.1
Role of AI	Hybrid human-in-the-loop	10	10.1
Behavioural direction	Reconfiguration	64	64.6
Behavioural direction	Attenuation	11	11.1
Behavioural direction	Amplification	4	4.0
Behavioural direction	Null	3	3.0
Behavioural direction	Contextual	17	17.2

Dimension	Category	n	% of 99
	(trust/adoption/intention)		

Note. Percentages use the full 99-study sample as the denominator. Behavioural direction was coded for 87 studies; 12 adoption/trust-focused studies were intentionally left without a directional classification. Broad method and AI-role groupings were created for descriptive presentation and do not replace the study-level classifications in the extraction matrix. Complete study-level details are provided in Supplementary Table S1.

5. Thematic Synthesis

The 99 included studies were organised into six themes: selective debiasing and portfolio discipline; behavioural amplification and possible reconfiguration; trust, resistance and calibrated reliance; human-AI interaction and social cues; investor heterogeneity and vulnerability; and information processing, delegation and GenAI cognitive partnership. The five reassessed studies contribute to the contextual trust/adoption strands, not to evidence of observed behavioural reconfiguration. The themes describe distinct mechanisms and levels of evidence rather than implying that every study measured investor-level behavioural effects.

5.1 Selective debiasing and portfolio discipline

Evidence-strength qualification (M089): In a cross-sectional questionnaire of 462 Indian individual investors, the authors of “Insight over instinct” report associations consistent with mitigation of anchoring and herding alongside persistence of overconfidence, representativeness and loss aversion. Because the design uses questionnaire responses and PLS-SEM rather than randomized AI exposure or observed pre/post trading, this study supports a bias-specific associational interpretation, not a causal estimate of AI debiasing. Its original extraction category remains frozen while the directional evidence is audited.

The strongest debiasing evidence comes from settings where automation constrains portfolio choices, rebalances portfolios or structures advice. D'Acunto, Prabhala and Rossi (2019) found improved diversification and reductions in the disposition effect and trend- or rank-related trading, while Rossi and Utkus (2024) reported greater indexation and international

diversification, lower fees and lower risk after robo-advice adoption. Lisauskiene, Daraskuviene and Butkus (2024) likewise found that robo-advisory interventions reduced willingness to continue holding losing stocks. These results suggest that automation is most effective when it removes discretion from bias-prone decision points rather than functioning as a generic rationality device.

Effects were nevertheless selective. Maheshwari and Samantaray (2026) reported weaker links between disposition/herding biases and irrational decisions in the presence of AI advisory services, whereas Bhatia et al. (2022), Wijaya, Liangga and Sutejo (2025) and Mulyono and Wiryakusuma (2025) reported null or bias-specific moderation. Saidalavi, Rasheed and Cholasseri (2025) found lower loss aversion but higher overconfidence and anchoring among AI-tool users. The relevant question is therefore which bias changes, through what mechanism and under what conditions.

5.2 Behavioural amplification and AI-induced reconfiguration

A second theme concerns behavioural costs created by interaction with AI. Ge (2026) found that around 30% of investors bypassed a trial stage before delegating, consistent with automation bias. The September update strengthens this interpretation: Xu and Liu (2026), using trading records, platform-usage logs and survey data, identify a U-shaped relationship between AI-assisted decision intensity and behavioural bias. Moderate use reduced bias and improved decision quality, whereas excessive reliance weakened independent judgment and increased automation bias.

Generative AI adds a metacognitive channel. Croom (2026) showed that interactive GenAI increased perceived ability and willingness to invest without equivalent improvement in actual information processing. Alongside evidence of greater retirement worry (Chhatwani, 2022) and simultaneous reductions and increases across different biases (Saidalavi, Rasheed and Cholasseri, 2025), these findings support behavioural reconfiguration: AI can alter the configuration of behavioural mechanisms rather than move investors uniformly towards or away from rationality.

5.3 Trust, resistance and calibrated reliance

Further source check (M035): Chen et al. (2025) surveyed 393 Chinese investors about robo-advisor-enabled green-investment adoption. The dependent endpoint was behavioural intention, with usefulness, ease of use and perceived trust in the explanatory model. The reported relationships should not be treated as observed ESG allocation or verified changes in investor bias. This distinction reinforces the separation between adoption mechanisms and realised investment decisions.

Verified source-level distinction (M060): Divekar et al. (2026) surveyed 282 existing robo-advisory users and modelled self-reported re-use intention. Their hindsight-bias measure predicts re-use intention ($\beta = 0.452$ in Table 6; $\beta = 0.432$ with controls in Table 10), alongside trust and satisfaction pathways. The study does not demonstrate that robo-advice changed hindsight bias or observed trading. It therefore informs post-adoption reliance intentions rather than establishing a measured behavioural-reconfiguration effect. Its original extraction label remains unchanged pending independent adjudication.

Trust was the most persistent mechanism, but greater trust was not uniformly beneficial. Cheng et al. (2019) linked supervisory control and trust-transfer relationships to robo-advisory trust; Guo and Waked (2026) identified trust as a central predictor of actual use; and Vijayananth and Saravanabhavan (2026) showed that algorithmic trust linked AI trading usage to investment behaviour and financial resilience, with human-AI interaction competency strengthening the trust-to-behaviour pathway. These studies establish trust as a transmission mechanism between system characteristics and reliance.

The behavioural evidence points to calibrated reliance rather than trust in isolation. Filiz et al. (2022) found delegation to an objectively superior robo-advisor in only about 40% of cases, illustrating under-reliance, while Ge (2026) identified premature delegation consistent with over-reliance. Advice-taking studies show that performance information, framing, perceived ability and integrity also shape reliance (Mahmud, Islam and Krishnan, 2025; Schütz, Schröder and Rennhak, 2023), whereas algorithm aversion and autonomy concerns can generate resistance (Goswami and Verma, 2026). Both insufficient and excessive reliance therefore carry behavioural costs.

5.4 Human-AI interaction, social cues and hybrid advice

Human-like characteristics and human involvement reshape investor responses independently of informational content. Conversationality and anthropomorphic cues can increase affective trust, perceived competence and recommendation acceptance (Hildebrand and Bergner, 2021; Plotkina, Orkut and Karageyim, 2024; Cao et al., 2026). Yet Back, Morana and Spann (2023) found that social design reduced advice seeking and partially reversed a robo-advisor's debiasing benefit, showing that social appeal and behavioural welfare can diverge.

Hybrid arrangements reinforce this point. Yang et al. (2025) found greater alignment with human-AI collaborative advice than with pure AI advice, particularly for riskier decisions, while Hermann and Alberhasky (2026) and Northey et al. (2022) show that human relationships and involvement can remain powerful. Social cues and human participation therefore alter acceptance through affect, competence and responsibility, but may also introduce new heuristics; the design objective should focus on appropriate use of anthropomorphism and human presence rather than maximising either feature.

5.5 Investor heterogeneity, literacy and vulnerability

Retail investors are heterogeneous. Financial capability, experience, age, gender, risk tolerance, personality and vulnerability repeatedly differentiated AI adoption and behavioural outcomes. Nazmi, Lye and Tay (2024) highlighted transparency and digital education among lower-income investors; Piehlmaier (2022) found that overconfidence increased robo-advice adoption while higher financial literacy reduced uptake; and Nabeshima and Kadoya (2026), Fan and Chatterjee (2020), Badgujjar, Manrai and Yadav (2026) and Cheong and Chung (2026) document further demographic, capability, personality and execution heterogeneity.

Financial literacy is especially conditional rather than uniformly protective. It can support effective use, reduce adoption or make investors more critical of conversational systems. Eichler and Schwab (2025) found subgroup differences in trust and perceived risk, while Xu and Liu (2026) show that financial literacy can shift the threshold at which AI assistance moves from bias correction towards excessive reliance. The synthesis therefore distinguishes objective financial literacy, AI/digital literacy, experience and subjective confidence rather than treating them as interchangeable.

5.6 Information processing, delegation and GenAI cognitive partnership

The sixth theme marks the transition from conventional robo-advice to interactive GenAI. Whereas robo-advisors typically convert investor inputs into recommendations or portfolios, GenAI can become part of information processing itself. Blankespoor, Croom and Grant (2026) show investors using GenAI to interpret, contextualise and simplify financial information, while Ecker et al. (2026) find that users adapt their queries across investment tasks and disclosure events. AI therefore shifts from an external adviser towards an interactive cognitive resource.

This cognitive role offers benefits alongside new vulnerabilities. Chang et al. (2026) associate ChatGPT diffusion with greater alignment between retail trading and AI-derived earnings-call information; Gong and Müller (2026) find that GPT-4 simplification improves accessibility and willingness to invest; and Hoang et al. (2025) links following ChatGPT advice to reduced cognitive dissonance in crypto trading. Yet Croom (2026) shows that perceived ability may rise faster than processing quality. GenAI can thus reduce information frictions while simultaneously encouraging cognitive offloading, dependence or an illusion of competence.

5.7 Integrative interpretation of the six themes

The six themes organise the frozen corpus but do not establish the prevalence of any behavioural direction. The five corrected studies specifically illustrate why intention and trust cannot be used as proxies for observed trading or changed bias. The proposed pathway from system design and investor characteristics through cognition and reliance to investment outcomes remains a testable conceptual synthesis; evidence for each link varies in design and strength. Interactive GenAI is distinguished where primary evidence supports a different decision-process mechanism.

Table 2. Six themes, mechanisms and representative evidence

Theme	Core mechanism	Typical behavioural direction	Representative studies	Synthesis
Selective debiasing and portfolio discipline	Automation, structured advice, rebalancing, bias-aware design	Mainly attenuation, but selective	D'Acunto et al. (2019); Rossi & Utkus (2024); Lissauskiene et al. (2024); Saidalavi et al. (2025)	AI can reduce specific biases and improve portfolio discipline, but does not eliminate behavioural bias generally.
Behavioural amplification and AI-induced reconfiguration	Automation bias, overconfidence, metacognitive misattribution, engagement effects	Amplification / reconfiguration	Ge (2026); Croom (2026); Chhatwani (2022); Saidalavi et al. (2025); Xu & Liu (2026)	AI can create new behavioural distortions or shift the form of existing ones.

Theme	Core mechanism	Typical behavioural direction	Representative studies	Synthesis
Trust, resistance and calibrated reliance	Initial trust, competence/integrity, privacy, aversion, advice weighting	Contextual mechanism; directional effect only if decision outcome measured	Cheng et al. (2019); Filiz et al. (2022); Mahmud et al. (2025); Schütz et al. (2023); Vijayananth & Saravanabhavan (2026)	The relevant objective is calibrated reliance rather than maximum trust.
		Contextual / conditional; outcome-specific when directly measured		
Human-AI interaction, social cues and hybrid advice	Anthropomorphism, conversationality, affective trust, human-in-loop		Hildebrand & Bergner (2021); Back et al. (2023); Yang et al. (2025); Cao et al. (2026)	Human-like design can improve acceptance but can also reduce advice seeking or introduce heuristic trust.
Investor heterogeneity, literacy and vulnerability	Capability, experience, demographics, personality, risk tolerance	Conditional	Piehlmaier (2022); Nourallah (2023); Nabeshima & Kadoya (2026); Cheong & Chung (2026); Xu & Liu (2026)	Investor characteristics materially change both adoption and behavioural outcomes.
Information processing, delegation and GenAI cognitive partnership	Simplification, querying, cognitive offloading, adaptive interaction	Attenuation / reconfiguration / amplification	Blankespoor et al. (2026); Ecker et al. (2026); Gong & Müller (2026); Croom (2026)	GenAI can act as a cognitive scaffold while also creating dependence or illusion-of-ability risks.

Note. Representative studies are illustrative rather than exhaustive. The full evidence base for each theme is retained in the Stage 6 coding and audit matrix.

Table 3. Key contradictions and their interpretive resolution

Contradiction	Evidence pattern	Interpretive resolution	Framework implication
Bias mitigation is selective	Some studies report reductions in disposition, herding or overconfidence; others report null effects or simultaneous amplification of other biases.	The same AI intervention can reduce one bias, leave another unchanged and intensify a third.	Model attenuation, amplification, reconfiguration and null effects as parallel possibilities.
Human-like design can help or backfire	Conversationality and anthropomorphism often increase trust/acceptance, while some social designs reduce advice seeking.	Humanisation alters affective trust and competence heuristics rather than guaranteeing better decisions.	Treat social/human cues as moderators of calibrated reliance.
Trust can be too low or too high	Algorithm aversion can prevent beneficial delegation; automation bias can lead to uncritical acceptance.	Trust is valuable only when it is calibrated to system quality and task demands.	Use calibrated reliance as the central mechanism.

Contradiction	Evidence pattern	Interpretive resolution	Framework implication
Financial literacy is not uniformly protective	Literacy can support effective use, reduce adoption, or make investors more critical of conversational AI.	Literacy changes interpretation and reliance rather than simply increasing acceptance.	Treat literacy as a boundary condition, distinct from confidence and AI literacy.
GenAI can support processing and inflate self-assessment	Interactive AI can simplify information and support judgment, while also increasing perceived ability without equivalent processing gains.	GenAI produces a dual cognitive effect: assistance plus metacognitive risk.	Distinguish the GenAI cognitive-partnership pathway from traditional robo-advice.

6. Integrative Conceptual Framework

The framework separates four linked layers: AI/platform and investor/context antecedents; cognitive and affective mechanisms; reliance and interaction states; and measured behavioural consequences and investment outcomes. The five author-reclassified adoption/trust studies inform the earlier mechanism and reliance layers, not the final behavioural-outcome layer. The framework is a synthesis of relationships proposed across heterogeneous studies, not a validated single causal pathway.

6.1 Framework architecture

Figure 2 presents the integrative conceptual framework derived from the thematic synthesis. The framework links AI/platform characteristics and investor/context characteristics to cognitive and affective mechanisms, reliance states, behavioural consequences and investment decision outcomes.

Figure 2. Integrative conceptual framework: how AI-assisted technologies influence retail-investor decision-making




Figure 2. Proposed integrative conceptual framework of AI-assisted retail-investor decision-making. Trust, adoption and stated intention are contextual mechanisms or intermediate states; they do not by themselves demonstrate a change in trading, allocation or behavioural bias. The visual rendering was created with OpenAI's image-generation tool and reviewed by the author; the conceptual structure and interpretation were developed by the author.

Figure 2 distinguishes system/investor antecedents, cognitive and affective mechanisms, reliance and interaction states, and measured behavioural consequences. Trust, perceived competence, social presence, adoption intention and willingness to rely on advice belong to intermediate or contextual layers. Only studies measuring an investor decision, allocation, trading outcome or directly assessed bias can support a claim about attenuation, amplification, reconfiguration or null behavioural effects. The five reclassified studies inform the contextual pathways but cannot be used as outcome-layer evidence. This boundary also distinguishes advice acceptance from subsequent execution and welfare.

Boundary conditions: financial literacy and AI literacy, investment experience, age, gender, personality, risk tolerance, vulnerability, task complexity, decision stakes, framing, market conditions and institutional context.

6.2 AI/platform and investor/context antecedents

The framework begins with AI/platform and investor/context characteristics. Investors do not respond to AI as a uniform technology: transparency, interactivity, anthropomorphism, personalisation, autonomy, human involvement and information quality alter trust and advice-taking (Back, Morana and Spann, 2023; Hildebrand and Bergner, 2021; Cao et al., 2026). Investor characteristics - including literacy, experience, knowledge, risk tolerance, age, gender, personality and vulnerability - condition these responses (Piehlmaier, 2022; Nabeshima and Kadoya, 2026; Cheong and Chung, 2026). System design and investor characteristics therefore form joint boundary conditions for subsequent mechanisms and outcomes.

6.3 Cognitive and affective mechanisms

The second layer captures cognitive and affective mechanisms. Trust is prominent, but perceived competence, risk, control, cognitive load, self-assessed ability, reassurance, anxiety and social presence also transmit system characteristics into behaviour. Experimental and field evidence links ability/integrity cues and human-AI presentation to trust and advice adherence (Schütz, Schröder and Rennhak, 2023; Yang et al., 2025). GenAI adds an information-processing mechanism: it can simplify and contextualise information while also increasing perceived ability without equivalent gains in processing quality (Blankespoor, Croom and Grant, 2026; Ecker et al., 2026; Gong and Müller, 2026; Croom, 2026). The framework therefore separates objective cognitive support from subjective self-assessment.

6.4 Reliance, delegation and interaction states

These mechanisms shape a continuum of reliance. Algorithm aversion can prevent use of objectively valuable advice (Filiz et al., 2022), whereas excessive or heuristic trust can produce premature delegation and automation bias (Ge, 2026). Between these extremes lies calibrated reliance, in which investors accept, challenge or override advice in proportion to system quality and task demands. The same layer includes delegation, human-AI collaboration, cognitive offloading and execution. Outcomes depend on advice quality and on whether advice is sought,

weighted, retained under human control and ultimately executed (Mahmud, Islam and Krishnan, 2025; Yang et al., 2025; Cheong and Chung, 2026).

6.5 Behavioural consequences: attenuation, amplification, reconfiguration and null effects

The final outcome layer distinguishes attenuation (reduced measured distortion), amplification (increased measured distortion), reconfiguration (a demonstrable change in the pattern of multiple behavioural mechanisms or decision stages) and null effects (no detected difference on the specified measured outcome). These categories are empirical questions, not consequences inferred from trust or willingness to use a platform. Evidence on simultaneous lower loss aversion and higher overconfidence/anchoring suggests candidate reconfiguration mechanisms, while design-related changes in advice seeking demonstrate that interaction effects can differ across endpoints. The framework therefore requires future studies to measure investor decisions and multiple biases, rather than using adoption or confidence as proxies for welfare.

6.6 From traditional robo-advice to interactive GenAI

Traditional robo-advice and interactive GenAI should not be treated as behaviourally equivalent. Robo-advisory systems usually follow a profile-algorithm-recommendation sequence, making trust, acceptance, delegation and portfolio discipline central. GenAI permits repeated dialogue, adaptive questioning, interpretation and cognitive offloading, so the investor may be thinking with AI rather than merely receiving advice. This cognitive-partnership pathway can improve information access while also creating metacognitive risk if users attribute AI-assisted performance to their own competence or become dependent without evaluating system limitations (Croom, 2026; Blankespoor, Croom and Grant, 2026; Ecker et al., 2026; Gong and Müller, 2026).

Accordingly, GenAI research should examine how interaction changes attention, confidence, learning, verification and perceived competence alongside whether users follow recommendations.

6.7 Evidence-grounded propositions

The framework yields eight propositions that summarise the main relationships identified in the synthesis. These propositions are intended as integrative statements and future empirical targets, not as claims that each pathway has already been tested in a single unified model.

Table 4. Evidence-grounded propositions derived from the conceptual framework

Proposition	Statement	Principal mechanism	Boundary conditions
P1 — Conditional behavioural effects	AI-enabled investment technologies can attenuate, amplify, reconfigure or leave unchanged behavioural biases rather than producing a uniform debiasing effect. Insufficient trust can produce algorithm aversion and under-use,	Bias-specific response to AI advice and design	Bias type; task; investor characteristics; AI role
P2 — Trust calibration	while excessive or heuristic trust can produce automation bias and over-reliance. Human-like, conversational and human-in-the-loop features reshape advice acceptance through affective trust, social presence and perceived competence, but these cues can also backfire. Investor characteristics systematically condition AI-assisted decision-making and behavioural outcomes.	Trust → reliance/delegation	Transparency; risk; experience; provider cues; system tenure
P3 — Human/social cue pathway	Interactive GenAI differs from traditional robo-advice because it enters the investor's information-processing process as a cognitive partner. Cognitive benefits from GenAI can coexist with metacognitive risk, including overestimation of ability and excessive reliance.	Affective trust; competence heuristics	Anthropomorphism; social role; transparency; involvement
P4 — Investor heterogeneity	High-quality algorithmic advice does not determine outcomes by itself; compliance, execution speed and control choices mediate realised benefits. Effective AI investment systems should be designed for calibrated reliance rather than maximum automation, trust or engagement.	Capability and vulnerability	Literacy; experience; age; gender; personality; risk tolerance
P5 — GenAI cognitive partnership		Adaptive information processing	Task complexity; interactivity; sophistication
P6 — GenAI dual effect		Cognitive support + misattribution/offloading	Interactivity; task complexity; user sophistication
P7 — Human execution matters		Advice → execution → outcome	Control; execution frictions; stakes; confidence
P8 — Calibrated AI design		Design for calibrated reliance	Literacy; vulnerability; decision stakes; institutional safeguards

6.8 Theoretical contribution of the framework

The theoretical contribution is a measurement-sensitive, conditional framework connecting adoption and trust mechanisms to reliance and ultimately to measurable investment outcomes. The five-case author reassessment demonstrates why contextual acceptance evidence must be kept separate from behavioural-effect claims. Attenuation, amplification, reconfiguration and null effects remain testable possibilities; their relative frequency is not established by the present descriptive coding counts. Interactive GenAI may introduce distinct cognitive and metacognitive pathways, but those mechanisms require direct tests of consequential decisions.

7. Discussion and Implications

The six themes and integrative framework indicate that AI-assisted investing is best understood as a behavioural interaction rather than a neutral technological intervention. The literature is moving from adoption questions towards post-adoption mechanisms - reliance, delegation, information processing and execution - because decision quality depends on system design, investor characteristics, context and how advice is interpreted and acted upon.

7.1 Theoretical implications

First, AI cannot be modelled as a uniform debiasing mechanism. Strong studies document improvements in diversification and selected biases, while others report null, selective or simultaneous positive and negative effects (D'Acunto, Prabhala and Rossi, 2019; Rossi and Utkus, 2024; Bhatia et al., 2022; Saidalavi, Rasheed and Cholasseri, 2025). Theory should therefore allow attenuation, amplification, reconfiguration and null effects. Second, trust should be modelled as calibration rather than a monotonic antecedent of use: Filiz et al. (2022) illustrate under-reliance, while Ge (2026) illustrates over-reliance. Algorithm aversion, calibrated reliance and automation bias are distinct states on a reliance continuum.

Third, social and affective design cues function as behavioural interventions rather than simple interface enhancements. Conversationality, anthropomorphism and hybrid arrangements alter perceived competence, affective trust and advice acceptance, yet can also discourage useful advice seeking (Back, Morana and Spann, 2023; Hildebrand and Bergner, 2021; Yang et al., 2025; Cao et al., 2026). Fourth, interactive GenAI introduces a cognitive mechanism beyond traditional robo-advice: it can support interpretation and adaptive search while inflating self-

assessed ability (Blankespoor, Croom and Grant, 2026; Ecker et al., 2026; Gong and Müller, 2026; Croom, 2026). Behavioural-finance models should therefore distinguish cognitive assistance from metacognitive accuracy.

7.2 Implications for retail investors

For retail investors, AI can strengthen portfolio discipline and reduce information-processing costs, but assistance should not be confused with independent competence. Effective education should distinguish financial literacy, AI literacy, investment experience and subjective confidence, and emphasise when outputs require challenge or independent verification. This is especially important for GenAI, where easier information processing may coexist with over-reliance or inflated self-assessment.

7.3 Implications for financial platforms and robo-advisory providers

For platforms and robo-advisory providers, the design objective should be calibrated reliance rather than maximum adoption, trust or engagement. Transparency, explanations, human support and personalisation are valuable only when they improve appropriate use rather than persuasion. Providers should make uncertainty and system limitations visible and evaluate behavioural outcomes such as advice override, turnover and monitoring alongside conventional engagement metrics.

7.4 Regulatory and investor-protection implications

For regulators, suitability and disclosure remain necessary but may be insufficient where AI shapes confidence, reliance and the framing of decisions. The strongest concern is investor heterogeneity: those who benefit most from low-cost guidance may also find it hardest to evaluate its limitations. Conversational GenAI further blurs the boundary between information provision and advice, increasing the importance of understandable uncertainty, auditability, safeguards and appropriate human review for high-stakes decisions.

7.5 Implications for the development of behavioural finance

For behavioural finance, AI shifts attention from isolated biases towards the architecture of human-technology decision-making. Established mechanisms such as loss aversion, overconfidence, herding and the disposition effect remain relevant, but AI can change when they

are activated, which become dominant and whether they are expressed through independent judgment or interaction with an automated system.

7.6 Integrative discussion

The evidence supports a conditional rather than deterministic interpretation of AI-assisted investing. Some studies measure changes in investment behaviour, whereas others examine trust, adoption or future-use intention. The five corrected records fall into the latter category. The framework integrates these different evidence levels without assuming that trust causes improved outcomes or that reconfiguration is the dominant verified effect. Its contribution is to identify testable links among system design, investor characteristics, cognition, reliance and consequential investment decisions.

Table 5. Summary of principal implications

Audience	Key implication	Behavioural risk to avoid	Practical direction
Theory	Move from adoption and uniform debiasing to conditional behavioural reconfiguration.	Treating trust or AI use as uniformly beneficial.	Model antecedents, mechanisms, reliance states and heterogeneous outcomes.
Retail investors	AI can support discipline and information processing, but assistance should not be confused with independent competence.	Overconfidence, automation bias and reduced verification.	Maintain critical review, understand uncertainty and distinguish financial knowledge from AI literacy.
Platforms/providers	Design for calibrated reliance rather than maximum trust or engagement.	Persuasive design that increases acceptance without improving welfare.	Use transparency, appropriate friction, personalisation and human support selectively.
Regulators	Investor protection should consider behavioural effects as well as disclosure and suitability.	Vulnerable users placing excessive reliance on opaque or conversational systems.	Require clear limitations, auditability, appropriate safeguards and stronger protection in high-stakes settings.
Behavioural finance	AI changes the architecture through which biases are expressed.	Treating technology as a passive channel.	Study human-technology interaction as part of the behavioural mechanism itself.

8. Future Research Agenda

The future research agenda follows directly from the review's unresolved tensions and is organised using a Theory-Context-Characteristics-Methodology (TCCM) logic. Six priority

streams move the field beyond cross-sectional adoption studies towards causal evidence on when AI improves decision quality, when it creates behavioural risks and which investors are most affected.

8.1 GenAI cognitive partnership and metacognitive risk

Interactive GenAI creates a distinct cognitive-partnership problem because investors can query, refine and repeatedly interact with the system. Emerging evidence indicates both informational benefits and metacognitive risk: GenAI can simplify and contextualise financial information, yet may increase perceived ability without equivalent processing gains (Blankespoor, Croom and Grant, 2026; Ecker et al., 2026; Gong and Müller, 2026; Croom, 2026). Future studies should compare no-AI, static-AI and interactive-GenAI conditions while jointly measuring comprehension, confidence, reliance, verification and longitudinal learning or dependence.

8.2 Calibrated trust and reliance

Trust research should move from asking whether trust increases adoption to identifying when reliance is calibrated to system quality. Repeated-interaction experiments can expose investors to correct advice, errors and explanation-based repair to track trust formation, deterioration and recovery, while distinguishing cognitive and affective trust, perceived competence, integrity, benevolence and institutional assurance. This would clarify when trust produces appropriate reliance, algorithm aversion or automation bias (Filiz et al., 2022; Ge, 2026).

8.3 Multi-bias behavioural reconfiguration

Behavioural reconfiguration requires multi-bias designs. Most studies examine one bias at a time, making it difficult to detect substitution across disposition effects, anchoring, herding, overconfidence and loss aversion. Repeated-measures experiments, longitudinal panels and consequential simulations should test several biases within the same investors and comparable AI conditions, directly asking whether reducing one distortion leaves another unchanged or makes another more influential.

8.4 Human-AI design and investor welfare

Human-like and hybrid designs should be evaluated against investor welfare rather than advice acceptance alone. Comparative experiments across pure AI, human-only and human-AI advice should measure portfolio quality, realised risk, diversification, compliance, turnover, execution

costs and performance while manipulating anthropomorphism, transparency and human involvement. This is necessary because social design may increase acceptance yet also weaken debiasing or encourage heuristic trust (Back, Morana and Spann, 2023; Yang et al., 2025).

8.5 Investor heterogeneity and vulnerability

Investor heterogeneity should be treated explicitly. Future work should separate financial literacy, AI/digital literacy, investment experience and subjective confidence, and use large field samples or preregistered subgroup analyses to identify who benefits, who is unaffected and who is harmed. Oversampling older adults, first-time investors, low-literacy households and financially constrained users would strengthen investor-protection relevance.

8.6 Execution, market context and real financial outcomes

Research should connect AI advice to realised outcomes. Administrative brokerage data, platform records and natural experiments can trace recommendation, acceptance, execution and portfolio consequences, including delay, override, transaction costs, diversification, turnover and performance. The same designs should examine market stress, when loss aversion, panic selling and anxiety may alter reliance. Cheong and Chung (2026) already show that compliance speed can determine whether investors realise the benefit of robo-advisor recommendations.

8.7 Methodological priorities

Methodologically, the field needs stronger behavioural and causal identification. Randomized experiments, longitudinal panels, natural experiments, administrative data and platform collaborations should complement cross-sectional surveys. Prompt logs, query histories, clickstreams, advice-weighting measures and execution records can reveal how reliance evolves, while mixed methods can explain why investors accept, reject or reinterpret algorithmic advice.

8.8 TCCM research agenda

Table 6 consolidates these priorities into a selective TCCM agenda linking each gap to a research question, recommended design, key variables and contribution.

Table 6. Priority TCCM future research agenda

Priority stream	TCCM gap	Future research question	Recommended design	Key variables	Contribution
GenAI cognitive	Theory /	How does	Controlled	Comprehension;	Establishes a

Priority stream	TCCM gap	Future research question	Recommended design	Key variables	Contribution
partnership	Characteristics	interactive GenAI alter investors' objective processing and metacognitive judgments compared with static AI assistance?	experiment + longitudinal interaction logs	self-assessed ability; reliance; query behaviour	GenAI-specific behavioural mechanism.
Calibrated trust	Theory / Methodology	When does trust produce appropriate reliance, algorithm aversion or automation bias?	Repeated-advice experiment with error and repair	Trust trajectory; advice weight; override; error sensitivity	Develops a non-linear theory of trust calibration.
Multi-bias reconfiguration	Characteristics / Methodology	When does AI attenuate one bias while amplifying or leaving another unchanged?	Multi-bias repeated-measures experiment	Disposition; herding; anchoring; overconfidence; loss aversion	Direct test of behavioural reconfiguration.
Human-AI welfare	Characteristics / Methodology	When do human-like or hybrid advisors improve investor welfare rather than merely increase acceptance?	Consequential field/factorial experiment	Compliance; portfolio quality; risk; performance; affective trust	Separates persuasion from welfare improvement.
Investor heterogeneity	Context / Characteristics	Which investor groups benefit from AI advice, which are unaffected and which are harmed?	Large field dataset + heterogeneous treatment-effect analysis	Literacy; AI literacy; age; gender; vulnerability; experience	Provides investor-protection segmentation.
Execution and market stress	Context / Methodology	How do execution frictions and market stress alter the behavioural and financial effects of AI advice?	Administrative brokerage data + natural experiment	Compliance delay; turnover; panic selling; risk; performance	Links AI advice to realised financial outcomes.

8.9 Integrative research direction

Taken together, the agenda shifts attention from whether investors adopt AI to the conditions under which AI changes judgment, reliance and financial outcomes. The strongest studies will observe real interaction, identify causal mechanisms, compare heterogeneous investors and connect behavioural responses to portfolio consequences. Testing AI design, investor characteristics and context jointly will show whether AI improves decisions, displaces existing biases, creates new vulnerabilities or produces different combinations of these effects across investors and market settings.

9. Limitations

This review has limitations that define the boundaries of its conclusions.

First, the formal search used Scopus and Business Source Ultimate. Web of Science Core Collection could not be searched through the available institutional access, so the review does not claim exhaustive coverage of every bibliographic database. Broad AI, investor and behavioural terms were used, while supplementary academic-search tools supported triangulation and gap checking without silently expanding the frozen sample.

Second, eligibility was limited to English-language, peer-reviewed journal articles from 2015-2026. This improved comparability but may omit non-English work, conference papers, working papers and other grey literature, particularly in fast-moving GenAI research. Earlier behavioural-finance studies were used for theory but entered the systematic sample only when they directly examined technology-assisted investing.

Third, eight reports could not be retrieved in full despite further access attempts and were recorded as not retrieved rather than excluded. The September update added Mohapatra et al. (2026) after institutional retrieval failed and an author copy was requested. The final synthesis therefore comprises the 99 studies fully assessed and retained at the 5 September 2026 cut-off, while acknowledging that a not-retrieved report could have met the criteria.

Fourth, screening, extraction and synthesis used a primarily single-reviewer workflow with documented consistency checks and manual thematic auditing, rather than independent dual review. Five disputed classifications were reassessed by the author against previously documented full-text endpoints and recoded as contextual; invited independent reviewers were unavailable, so this must not be described as independent adjudication. The original codes and five-case sensitivity comparison are retained to make the decision auditable. Remaining directional records have not all received full source-level validation, limiting prevalence and causal interpretations.

Fifth, the evidence base is methodologically heterogeneous. Cross-sectional surveys and structural models dominate, alongside a smaller group of experiments, field studies, transaction records and administrative datasets. Evidence-strength ratings were therefore used to temper causal claims, but synthesis necessarily combines studies with different inferential strength.

Sixth, technology labels are fluid. Robo-advisor, AI-enabled advice, algorithmic advice, conversational AI and GenAI can describe systems with substantially different capabilities. Coding therefore focused on the role AI played in the investor decision process rather than terminology alone.

Seventh, the literature is geographically uneven, with substantial evidence from India, the United States, China and other Asian markets. Differences in regulation, culture, financial literacy and access to advice limit uniform generalisation across retail-investor populations.

Finally, GenAI is evolving faster than the academic publication cycle. The Scopus and Business Source Ultimate searches were therefore updated immediately before submission on 5 September 2026 using the same eligibility logic, adding two eligible studies. The framework should be read as an evidence-based synthesis at that cut-off rather than a final model of AI-assisted investor behaviour.

Future reviews would benefit from broader database/language coverage, independent dual screening, clearer reporting of AI-system characteristics and greater use of longitudinal, field and transaction data.

10. Conclusion

Artificial intelligence is changing retail investment decision-making in ways that cannot be reduced to rational algorithms versus biased investors. Across 99 peer-reviewed studies, AI can improve portfolio discipline and reduce selected distortions, but these benefits are neither universal nor behaviourally neutral.

The review proposes a measurement-sensitive framework for conditional AI-assisted retail-investor decision-making. The documented five-case author reassessment changes the descriptive coding distribution to 64 reconfiguration, 11 attenuation, four amplification, three null and 17 contextual records, while preserving $n=99$. Neither the original nor revised counts establish the prevalence of verified causal behavioural effects; independent second-reviewer adjudication was unavailable, and further source-level validation is needed for other directional records. The framework should therefore be interpreted as a structured research agenda, not evidence that reconfiguration predominates empirically.

Trust is central but should not be maximised. Under-reliance can prevent investors from benefiting from useful advice, whereas excessive or heuristic trust can encourage automation bias and uncritical delegation. The more appropriate objective is calibrated reliance: weighting AI advice in a manner consistent with system quality, task uncertainty and the investor's ability to evaluate the recommendation.

The review also distinguishes conventional robo-advice from interactive GenAI. Robo-advisors mainly influence profiling, portfolio construction and rebalancing, while GenAI can enter the reasoning process through interpretation, questioning and iterative judgment. This cognitive partnership can reduce information frictions while creating metacognitive risk when easier task performance is mistaken for greater independent competence.

Behavioural-finance theory should therefore move beyond adoption and uniform debiasing assumptions. Investors need financial and AI literacy. Providers should design for appropriate reliance rather than maximum engagement, while regulators should consider how adaptive systems shape confidence, decision processes and the final recommendation.

AI is unlikely to remove behavioural finance from investing; it changes its form. The key question is how human and machine judgment interact, where behavioural vulnerabilities migrate and when that interaction improves investor welfare. The framework and research agenda provide a basis for testing these questions as the technology and evidence evolve.

Data availability statement

The review dataset supporting this article is provided in Supplementary Table S1, which lists all included studies and core study characteristics. Additional screening and coding audit records are available from the corresponding author on reasonable request.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the author used OpenAI ChatGPT to assist with manuscript organisation, consistency checking, language refinement and readability. After using this tool,

the author independently reviewed, verified and edited the content as needed and takes full responsibility for the content of the publication.

References

- Akhtar, F., Akhtar, S., & Laeeq, M. (2025). Evolution of robo-advisors: A literature review and future research agenda. *International Journal of Consumer Studies*, 49(6), e70131. <https://doi.org/10.1111/ijcs.70131>
- Ashwini C G, & Ahmed, K. T. (2026). AI-enabled FinTech and investment decision-making: An integrative review of capabilities, behavioural adoption, and governance. *International Journal of Economic Practices and Theories*, 2026, 6258–6272. <https://doi.org/10.52710/ijept.891>
- Back, C., Morana, S., & Spann, M. (2023). When do robo-advisors make us better investors? The impact of social design elements on investor behavior. *Journal of Behavioral and Experimental Economics*, 103, 101984. <https://doi.org/10.1016/j.socec.2023.101984>
- Badgujar, K., Manrai, R., & Yadav, P. D. (2026). Understanding the factors affecting robo advisory adoption intention under the lens of Big Five personality and FOMO using the SOR model. *Vision: The Journal of Business Perspective*. Advance online publication. <https://doi.org/10.1177/09722629261475030>
- Bhatia, A., Chandani, A., Divekar, R., Mehta, M., & Vijay, N. (2022). Digital innovation in wealth management landscape: the moderating role of robo advisors in behavioural biases and investment decision-making. *International Journal of Innovation Science*, 14(3/4), 693–712. <https://doi.org/10.1108/IJIS-10-2020-0245>
- Blankespoor, E., Croom, J., & Grant, S.M. (2026). Generative AI and investor processing of financial information. *Journal of Accounting and Economics*, Advance online publication, 101908. <https://doi.org/10.1016/j.jacceco.2026.101908>
- Cao, Z., Yu, K., Guo, L., & Zhang, Y. (2026). Round versus square: Exploring the influence of facial shape in robo-advisors on consumer acceptance of investment advice. *Information & Management*, 63(2), 104291. <https://doi.org/10.1016/j.im.2025.104291>
- Cardillo, G. & Chiappini, H. (2024). Robo-advisors: A systematic literature review. *Finance Research Letters*, 62, 105119. <https://doi.org/10.1016/j.frl.2024.105119>
- Chang, A.Y., Dong, X., Martin, X., & Zhou, C. (2026). AI democratization and trading inequality. *Journal of Accounting Research*, 64(3), 1287–1331. <https://doi.org/10.1111/1475-679x.70063>
- Cheng, X., Guo, F., Chen, J., Li, K., Zhang, Y., & Gao, P. (2019). Exploring the trust influencing mechanism of robo-advisor service: A mixed method approach. *Sustainability*, 11(18), 4917. <https://doi.org/10.3390/su11184917>
- Cheong, S. C. & Chung, J. H. (2026). Boys will be boys, but robots can help them: Gender, algorithmic compliance, and portfolio performance. *Economics Letters*, 262, 112874. <https://doi.org/10.1016/j.econlet.2026.112874>
- Chhatwani, M. (2022). Does robo-advisory increase retirement worry? A causal explanation. *Managerial Finance*, 48(4), 611–628. <https://doi.org/10.1108/MF-05-2021-0195>
- Croom, J. (2026). Interactivity and illusions of ability: How using generative AI affects investor judgments. *Journal of Accounting Research*, 64(2), 681–719. <https://doi.org/10.1111/1475-679X.70017>
- 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. <https://doi.org/10.1093/rfs/hhz014>
- 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>
- 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*, 64(3), 1375–1426. <https://doi.org/10.1111/1475-679x.70051>
- Eichler, K. S. & Schwab, E. (2025). Differences in investment behavior – A quantitative study on the effects of robo-advisors and human financial advisors. *Journal of Applied Business and Economics*, 27(2). <https://doi.org/10.33423/jabe.v27i2.7607>
- Fan, L. & Chatterjee, S. (2020). The utilization of robo-advisors by individual investors: An analysis using diffusion of innovation and information search frameworks. *Journal of Financial Counseling and Planning*, 31(1), 130–145. <https://doi.org/10.1891/JFCP-18-00078>
- Filiz, I., Judek, J.R., Lorenz, M., & Spiwoks, M. (2022). Algorithm aversion as an obstacle in the establishment of robo advisors. *Journal of Risk and Financial Management*, 15(8), 353. <https://doi.org/10.3390/jrfm15080353>

- Ge, R. (2026). Skipping the trial: An empirical investigation of automation bias in robo-advisor adoption. *Computers in Human Behavior Reports*, 23, 101158. <https://doi.org/10.1016/j.chbr.2026.101158>
- Gong, Z. & Müller, S. (2026). No matter your financial literacy: Simplicity wins when choosing a fund. *Journal of Behavioral and Experimental Finance*, 50, 101188. <https://doi.org/10.1016/j.jbef.2026.101188>
- Goswami, D. & Verma, B. (2026). Why investors resist: A psychological perspective on adoption resistance to financial robo-advisors. *Journal of Internet Commerce*, 25(1), 38–69. <https://doi.org/10.1080/15332861.2025.2566660>
- Guo, Y. & Waked, H.N. (2026). Trust-mediated adoption of AI robo-advisors in inland China: An extended UTAUT perspective. *Journal of Financial Services Marketing*, 31, 11. <https://doi.org/10.1057/s41264-025-00339-2>
- Hermann, E. & Alberhasky, M. (2026). The trusted partner for financial decision making: Romantic partner or AI? *Computers in Human Behavior*, 177, 108877. <https://doi.org/10.1016/j.chb.2025.108877>
- Hildebrand, C. & Bergner, A. (2021). Conversational robo advisors as surrogates of trust: onboarding experience, firm perception, and consumer financial decision making. *Journal of the Academy of Marketing Science*, 49(4), 659–676. <https://doi.org/10.1007/s11747-020-00753-z>
- Hoang, S. D., Nguyen, T. H.-H., Dey, S. K., & Thi Thu, H. D. (2025). Beyond the hype: AI advice and investor dissonance in crypto trading. *Current Psychology*, 44(5), 3313–3325. <https://doi.org/10.1007/s12144-025-07430-w>
- Kahneman, D. & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–292. <https://doi.org/10.2307/1914185>
- Lisauskiene, N. & Darskuvienė, V. (2021). Linking the robo-advisors phenomenon and behavioural biases in investment management: An interdisciplinary literature review and research agenda. *Organizations and Markets in Emerging Economies*, 12(2), 459–477. <https://doi.org/10.15388/omee.2021.12.65>
- Lisauskiene, N., Darskuvienė, V., & Butkus, M. (2024). Passive vs active robo-advisors and disposition effect. Moderating role of gender and financial literacy. *Baltic Journal of Economics*, 24(2), 239–260. <https://doi.org/10.1080/1406099X.2024.2422673>
- 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>
- Maheshwari, H. & Samantaray, A.K. (2026). Navigating investment challenges: financial literacy and digital advisory services in countering behavioural biases. *International Journal of Sociology and Social Policy*, 46(3/4), 409–425. <https://doi.org/10.1108/IJSSP-12-2024-0620>
- Mahmud, H., Islam, N., & Krishnan, S. (2025). Human–robo-advisor collaboration in decision-making: Evidence from a multiphase mixed methods experimental study. *Decision Support Systems*, 198, 114541. <https://doi.org/10.1016/j.dss.2025.114541>
- Mohapatra, A. K., Sahoo, A. P., Tripathy, P., Matta, R., & Saxena, A. (2026). Adoption of explainable artificial intelligence in retail investors' decision-making: Evidence from India. *Indian Journal of Finance*, 20(1), 24–37. <https://doi.org/10.17010/ijf/2026/v20i1/175903>
- Mulyono, N. K. & Wiryakusuma, I. B. Y. (2025). Artificial intelligence as a game-changer in investment decisions: A financial literacy perspective. *JMBI UNSRAT: Jurnal Ilmiah Manajemen Bisnis dan Inovasi Universitas Sam Ratulangi*, 12(3), 1281–1292. <https://doi.org/10.35794/jmbi.v12i3.64363>
- Nabeshima, H. & Kadoya, Y. (2026). Investor heterogeneity in preferences for AI-based financial services: Evidence from Japanese online investors. *Risks*, 14(7), 154. <https://doi.org/10.3390/risks14070154>
- Nazmi, A.N.A., Lye, C.-T., & Tay, L.-Y. (2024). Promoting robo-advisor adoption among B40 in Malaysia through advisory transparency and UTAUT models. *Engineering, Technology and Applied Science Research*, 14(6), 18727–18733. <https://doi.org/10.48084/etasr.8289>
- Northey, G., Hunter, V., Mulcahy, R., Choong, K., & Mehmet, M. (2022). Man vs machine: how artificial intelligence in banking influences consumer belief in financial advice. *International Journal of Bank Marketing*, 40(6), 1182–1199. <https://doi.org/10.1108/IJBM-09-2021-0439>
- Nourallah, M. (2023). One size does not fit all: Young retail investors' initial trust in financial robo-advisors. *Journal of Business Research*, 156, 113470. <https://doi.org/10.1016/j.jbusres.2022.113470>
- Nourallah, M., Öhman, P., Walther, T., & Nguyen, D. K. (2026). Financial robo-advisors: A scoping review and future research directions. *Journal of Behavioral and Experimental Finance*, 49, 101158. <https://doi.org/10.1016/j.jbef.2026.101158>
- Odean, T. (1998). Are investors reluctant to realize their losses? *The Journal of Finance*, 53(5), 1775–1798. <https://doi.org/10.1111/0022-1082.00072>

- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., McGuinness, L. A., Stewart, L. A., Thomas, J., Tricco, A. C., Welch, V. A., Whiting, P., & Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
- Piehlmaier, D.M. (2022). Overconfidence and the adoption of robo-advice: why overconfident investors drive the expansion of automated financial advice. *Financial Innovation*, 8, 14. <https://doi.org/10.1186/s40854-021-00324-3>
- Plotkina, D., Orkut, H., & Karageyim, M.A. (2024). Give me a human! How anthropomorphism and robot gender affect trust in financial robo-advisory services. *Asia Pacific Journal of Marketing and Logistics*, 36(10), 2689–2705. <https://doi.org/10.1108/APJML-09-2023-0939>
- Rossi, A.G. & Utkus, S. (2024). The diversification and welfare effects of robo-advising. *Journal of Financial Economics*, 157, 103869. <https://doi.org/10.1016/j.jfineco.2024.103869>
- Saidalavi, K., Rasheed, K., & Cholasseri, A. (2025). Influence of AI Trading Tools on Cognitive Biases and Investor Behavior: An Exploratory Study. *IUP Journal of Accounting Research & Audit Practices*, 24(3), 562–580. <https://doi.org/10.71329/IUPJARAP/2025.24.3.562-580>
- Santosa, K. N., Nadlifatin, R., & Subriadi, A. P. (2026). A Systematic Review on Digital Investment Platform Adoption Through the Lens of Technology, Risk, Trust, Quality, and User Experience Dimensions. *Journal of Information Systems Engineering and Business Intelligence*, 12(2), 236–252. <https://doi.org/10.20473/jisebi.12.2.236-252>
- Schütz, T., Schröder, C., & Rennhak, C. (2023). Acceptance of automated investment advisory: An experimental study of the relevance of trust attributes of a robo-advisor. *Die Unternehmung: Swiss Journal of Business Research and Practice*, 77(2), 185–201. <https://doi.org/10.5771/0042-059x-2023-2-185>
- Shefrin, H. & Statman, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of Finance*, 40(3), 777–790. <https://doi.org/10.1111/j.1540-6261.1985.tb05002.x>
- Vijayananth, M., & Saravanabhavan, N. (2026). Beyond algorithmic trust: human-AI interaction competency strengthens AI-driven financial decision-making and financial resilience. *Frontiers in Artificial Intelligence*, 9, 1867266. <https://doi.org/10.3389/frai.2026.1867266>
- Wijaya, L. I., Liangga, V. M., & Sutejo, B. S. (2025). Do robo-advisors as moderating variables weaken the overconfidence and loss aversion behavior bias of young investors' mutual fund investment decisions? *Jurnal Aplikasi Bisnis dan Manajemen*, 11(2), 349–360. <https://doi.org/10.17358/jabm.11.2.349>
- Xu, Z., & Liu, W. (2026). Nonlinear effects of machine learning-assisted investment decisions on investor behavior and asset pricing efficiency. *Mathematics*, 14(15), 2683. <https://doi.org/10.3390/math14152683>
- Yang, C., Bauer, K., Li, X., & Hinz, O. (2025). My advisor, her AI, and me: Evidence from a field experiment on human–AI collaboration and investment decisions. *Management Science*, 72(1), 242–264. <https://doi.org/10.1287/mnsc.2022.03918>
- Meng, K., Mahapatra, M. S., & Xiao, J. J. (2025). Artificial intelligence and consumer financial behavior: A systematic literature review and agenda for future research. *Journal of Consumer Behaviour*, 24(4), 1755–1786. <https://doi.org/10.1002/cb.2497>