

Article

Financial Digital Twins and Conversational AI in Robo-Advisory: Evidence from a Scenario-Based Randomized Experiment

Marco I. Bonelli 

Venice School of Management, Ca' Foscari University of Venice, San Giobbe, Cannaregio 873, Fondamenta San Giobbe, 30121 Venice, Italy; mibonelli6@gmail.com

Abstract

Robo-advisors have expanded access to automated investment services, but many platforms continue to rely on relatively static onboarding procedures and limited forms of user interaction. This study examines how participants with investment experience respond to two next-generation robo-advisory design features: financial digital twins, understood as dynamic investor profiles that integrate goals, risk tolerance, cash-flow patterns, and anticipated life events, and conversational artificial intelligence (AI), understood as an interactive interface for explaining recommendations. Using a scenario-based randomized 2×2 online experiment, 336 adult respondents with self-reported investment experience, recruited through professional and academic networks, were assigned to one of four robo-advisor scenarios that varied the personalization architecture, standard profile versus digital twin, and the interface style, plain dashboard versus conversational AI, while holding the portfolio recommendation constant. The results show that digital-twin personalization increases perceived personalization and privacy concern, indicating that more adaptive advisory architectures may be viewed as both more relevant and more data-intensive. Conversational AI increases the perceived interactive quality of the advisory experience, while selected willingness-related patterns, especially in the combined digital-twin and conversational-AI condition, are treated as exploratory because several secondary composites displayed limited internal consistency. The strongest confirmatory emphasis is therefore placed on perceived personalization and privacy concern, and the remaining findings are best interpreted as scenario-based investor responses rather than evidence of actual adoption behavior or confirmed psychological mechanisms. The study contributes to behavioral FinTech research by clarifying the personalization–privacy tension in AI-enabled robo-advisory services and by offering design implications for more transparent, interactive, and responsibly personalized digital wealth-management systems.



Academic Editors: Stefanos Balaskas,
Yannis C. Stamatou, George S.
Androulakis and Roberto Vergallo

Received: 22 April 2026

Revised: 15 June 2026

Accepted: 22 June 2026

Published: 1 July 2026

Copyright: © 2026 by the author.
Licensee MDPI, Basel, Switzerland.
This article is an open access article
distributed under the terms and
conditions of the [Creative Commons
Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: robo-advisors; financial digital twins; conversational AI; personalization; privacy concern; FinTech

JEL Classification: G11; G24; G41; O33; C91

1. Introduction

Robo-advisory has become one of the most visible forms of automation in retail financial services because it lowers the cost of investment advice while changing how investors evaluate and delegate financial judgment. By converting investor information into risk profiles, portfolio recommendations, rebalancing rules, and monitoring procedures, robo-advisors replace key elements of the traditional advisory relationship with algorithmic

processes that users must assess without the reassurance of direct human interaction [1–4]. This makes robo-advisory a central setting for behavioral FinTech research, where adoption depends not only on technical accuracy or portfolio efficiency, but also on whether the system appears useful, trustworthy, understandable, privacy-respecting, and aligned with the investor's circumstances.

Much of the early robo-advisory model still relies on standardized onboarding questionnaires that capture a narrow set of variables, including age, income, investment horizon, objectives, and risk tolerance. Although this structure supports scalability and consistency, it can reduce the investor to a static profile in a context where financial advice depends on changing circumstances, perceived suitability, and the ability to explain why a recommendation fits. Investors may therefore evaluate automated advice not only by the portfolio proposed, but by the architecture through which the system represents them and communicates its recommendation. The next stage of robo-advisory research should consequently move beyond general adoption and examine how specific advisory designs shape perceptions of personalization, privacy, interaction quality, and willingness to rely on automated advice.

Financial digital twins and conversational AI sharpen this design question because they change both how the investor is represented and how the recommendation is explained. A financial digital twin is a dynamic, permission-based representation of the investor's financial life, potentially incorporating goals, income patterns, spending behavior, liabilities, account structure, household circumstances, tax considerations, and expected life events [5–7]. Conversational AI instead operates at the interface level, allowing advice to be communicated through a more interactive and responsive explanation rather than a static dashboard [7,8]. Although both features promise deeper personalization and clearer communication, their behavioral effects are theoretically ambiguous: a richer investor model can make advice feel more relevant while exposing broader data requirements, whereas a conversational interface can improve comprehension while increasing social presence and reliance on automated recommendations.

Theoretical understanding of this tension requires bringing together three research streams that are often treated separately in robo-advisory studies: technology adoption, AI trust, and privacy calculus. TAM and UTAUT emphasize perceived usefulness, ease of use, and facilitating conditions as drivers of acceptance [9,10], while trust theory and AI-trust research stress that users must judge the competence, reliability, transparency, and integrity of systems that perform consequential tasks on their behalf [11,12]. Privacy-calculus research adds that users weighed expected benefits against the costs and risks of disclosing personal data [13]. These mechanisms converge in robo-advisory because personalization may require information about income, assets, liabilities, spending patterns, tax position, household needs, and future plans, while the resulting advice affects financial risk and long-term wealth allocation. The central theoretical issue is therefore whether investors interpret richer personalization as a benefit, a privacy cost, or both simultaneously.

Prior research shows that robo-advisor acceptance is shaped by trust, perceived usefulness, financial literacy, digital-finance familiarity, provider credibility, perceived expertise, and preferences for human or hybrid advice [2–4,14–18]. Related studies show that personalization can influence trust in recommendation agents [19], conversational robo-advisors can affect onboarding and financial decisions [8], and social design elements can shape investor behavior [20]. Conceptual work has also identified financial digital twins and natural-language technologies as promising directions for adaptive robo-advisory systems [5–7]. Yet limited evidence shows how investors respond when these two features are experimentally varied within the same advisory setting while the portfolio recommendation is held constant. The literature therefore still lacks clear evidence on whether digital-twin

personalization increases perceived relevance, privacy concern, or both, and whether conversational AI improves the perceived interactional quality of automated advice.

To address this gap, the study uses a scenario-based randomized 2×2 online experiment with 336 adult respondents who reported investment experience. Participants were assigned to one of four robo-advisory scenarios that varied the personalization architecture, standard profile versus financial digital twin, and the interface style, plain dashboard versus conversational AI, while the recommended portfolio remained constant across conditions. This design isolates the behavioral effect of advisory architecture from differences in financial advice content. The analysis examines whether digital-twin personalization increases perceived personalization and privacy concern, whether conversational AI increases perceived social presence, and whether these features are associated with exploratory indicators of stated willingness, including trust-related responses, adoption intention, advice compliance, portfolio acceptance, and hypothetical allocation willingness.

Within these design boundaries, the paper does not claim to observe actual adoption behavior in a live robo-advisory platform or to validate a complete psychological model of robo-advisor adoption. It provides controlled experimental evidence on how investors evaluate AI-enabled advisory scenarios that differ in data architecture and communication interface. The measurement diagnostics further require a clear hierarchy of interpretation: perceived personalization and privacy concern receive the strongest inferential emphasis, social presence is interpreted cautiously, and downstream indicators such as trust, usefulness, adoption intention, advice compliance, fiduciary quality, data control, and data-sharing willingness are treated as exploratory response patterns rather than validated latent constructs.

The results show that digital-twin personalization increases both perceived personalization and privacy concern, suggesting that richer advisory architectures are experienced as more relevant but also more data-intensive. Conversational AI increases perceived social presence, indicating that an interactive interface can alter the advisory experience even when the portfolio recommendation is unchanged. Exploratory patterns further suggest that the combined digital-twin and conversational-AI condition is associated with stronger selected willingness-related responses, especially hypothetical allocation willingness, although these findings are interpreted cautiously.

The main contribution is therefore not the broad claim that these technologies automatically increase robo-advisor adoption, but the more specific insight that AI-enabled advisory design creates a personalization–privacy trade-off, and that the explanation interface may shape how investors experience it. The contribution is therefore design-theoretical rather than merely technological: it shows how the data architecture of personalization and the interaction architecture of explanation jointly shape investor perceptions under controlled advisory-content conditions.

The study contributes to behavioral FinTech and robo-advisory research by shifting attention from general acceptance to the design architecture of automated advice, showing that the same portfolio recommendation can be evaluated differently depending on how the investor is represented and how the advice is communicated. It brings privacy calculus into the robo-advisory setting by showing that deeper personalization can increase perceived relevance while also heightening concern about data use, and it frames conversational AI as an interactional layer whose value should be judged not only by engagement or willingness, but by explanation, transparency, and responsible investor understanding.

These findings also carry practical implications for digital wealth-management design: financial digital twins should be auditable and permission-based, while conversational AI should function as an explanation layer rather than as a persuasive interface that merely increases comfort with automated recommendations.

2. Literature Review and Hypotheses Development

2.1. Robo-Advisors and AI-Driven Financial Advice

Robo-advisors are digital financial-advice platforms that collect investor information, assess risk tolerance, recommend portfolios, and automate investment management. Their appeal lies in scalability, relatively low fees, automated rebalancing, broader access to financial advice, and a more disciplined investment process for retail users who may not otherwise rely on traditional advisors [1–4].

At the same time, robo-advisory raises a persistent challenge because investors are asked to rely on algorithmic recommendations in decisions involving personal wealth, risk exposure, and long-term financial goals, often without the relational reassurance or contextual interpretation provided by human advisors. For this reason, adoption depends not only on functionality, but also on whether the system is perceived as useful, understandable, credible, and aligned with investor interests [2–4]. Technology-adoption research provides the basic framework for explaining these responses.

The Technology Acceptance Model emphasizes perceived usefulness and ease of use [9], while UTAUT broadens the explanation by incorporating social influence and facilitating conditions [10]. In financial services, however, these models are informative but incomplete, because they do not fully capture the trust, expertise, and risk judgments that accompany delegated investment decisions. A meta-analytical study confirms this broader view, showing that trust, attitude, perceived usefulness, and ease of use remain central drivers of robo-advisor adoption, while also indicating that institutional and cultural context can condition these effects [21]. The implication is that the next stage of research should move beyond general acceptance and examine how particular design architectures shape investor responses.

2.2. Trust, Expertise, and Adoption in Automated Financial Advice

Trust occupies a central place in robo-advisory because users delegate part of their financial judgment to a non-human system. General trust theory defines trustworthiness through ability, benevolence, and integrity [11], and AI-trust research extends this logic by emphasizing competence, reliability, transparency, and the cues embedded in human-machine interaction [12]. In the robo-advisory context, these issues become even more important because users are not only evaluating an algorithm, but also evaluating the firm behind it, the broader advisory environment, and the expected consequences of relying on automated advice.

Recent studies reinforce this multi-layered understanding of trust in automated financial advice. Nourallah et al. show that young retail investors build initial trust in financial robo-advisors through perceived benefits, risk evaluation, and confidence in the system [15], while Cao et al. distinguish between technology trust, firm-specific trust, and system trust, showing that investors evaluate not only the tool itself but also the provider and the wider financial setting in which advice is embedded [16]. Lourenço et al. similarly show that the type of firm offering robo-advice shapes perceptions of trust and expertise [14], and Luo et al. add that provider reputation, fiduciary signals, customer education, privacy protection, and trust contribute to willingness to use robo-advisors and allocate assets through them [17].

Evidence from Asian and emerging-market settings is especially relevant because robo-advisory adoption may also depend on financial literacy, institutional trust, privacy protection, and the continuing role of human advice. Evidence from India shows that trust, performance expectancy, and preferences for human advice remain important in explaining robo-advisor adoption [18], while recent Asian digital-finance studies suggest that socially present AI interfaces and chatbot-mediated recommendation processes can

shape user experience, trust-related evaluations, and willingness to continue relying on automated advice [22,23]. What emerges from this literature is a consistent point: robo-advisor adoption is not explained by usability alone, but by a broader judgment about competence, credibility, institutional context, and the acceptability of delegating financial decisions to an algorithmic system.

2.3. Financial Digital Twins in Robo-Advisory Services

Financial digital twins extend this debate because they promise a richer and more adaptive form of personalization than conventional robo-advisory profiles. In this context, a financial digital twin can be understood as a dynamic, permission-based representation of an investor's financial life, incorporating risk tolerance, financial goals, income patterns, spending behavior, liabilities, account structure, household needs, tax context, and expected life events. Rather than relying on a one-time onboarding questionnaire, a digital-twin system can update the investor profile as circumstances change and, in principle, generate more context-sensitive recommendations.

Anshari et al. identify digital twins as a potential next frontier for robo-advisory because they may enable more customized and adaptive financial management [5]. Bhat and Krishnan place digital twins within a wider transformation of financial systems, emphasizing their role in simulation, personalization, and decision support [24], while earlier work on robo-advisory innovation also highlights digital-twin integration as a plausible strategic path for improving personalization and service quality [6,7]. Yet this line of work remains largely conceptual, the causal evidence is still limited, and it is not yet clear whether investors respond more favorably to a robo-advisor that presents itself as a dynamic financial twin than to one based on a standard risk-profile questionnaire.

The likely attraction of digital twins lies in their promise of deeper personalization because recommendation-agent research has long shown that personalization and familiarity can strengthen trust and adoption when users perceive the advice as more relevant and user-centered [19]. Recent evidence from Germany also shows that investor profiling and portfolio recommendations remain heterogeneous across robo-advisory platforms, which keeps questions of suitability and transparency open [25].

In the investment setting, that mechanism should be especially important, since suitability depends on understanding the investor's goals, constraints, and life circumstances rather than a narrow set of profile variables. At the same time, a more intensive personalization architecture introduces a competing mechanism, because broader personalization requires broader data access.

Privacy-calculus research suggests that users weigh expected benefits against privacy costs [13], and in the robo-advisory setting, those costs may be particularly salient when data extend to income, spending, tax, household, and account information. Ethical research on AI-based financial advice further reinforces this tension by highlighting concerns around explainability, vulnerability, responsibility, and the possibility that users may overtrust systems whose operations they do not fully understand [26]. Digital-twin personalization may therefore increase perceived usefulness and trust, while also increasing privacy concern.

2.4. Conversational AI and Social Presence in Financial Advice

Conversational AI addresses a different limitation of conventional robo-advisors, namely the weak interaction that characterizes many dashboard-based platforms. A static interface may present a portfolio allocation and a brief explanation, but it does not necessarily help the investor understand why the recommendation is suitable, what assumptions underpin it, or how it relates to personal goals. Conversational systems can make the advi-

sory experience feel more responsive by explaining the rationale of the recommendation, answering likely follow-up questions, and presenting the advice in a more interactive way.

The literature increasingly suggests that this relational layer matters, a point illustrated by Hildebrand and Bergner, who show that conversational robo-advisors can act as surrogates of trust by improving onboarding experience, firm perception, and consumer financial decision-making [8]. Tang and Son further show that the sociality of AI digital human advisors can increase user experience value through social presence in digital finance [22], while Wang and Aw show that recommendation adoption and stickiness to the robo-advisor are shaped by trust, perceptions of the bot, and the broader interaction with the advisory system [23].

These studies indicate that conversational features are not merely cosmetic but can influence how investors interpret the advice, how much confidence they place in the system, and how willing they are to continue relying on it over time. At the same time, the literature does not present conversational design as an unqualified benefit. Back et al. show that social design elements in robo-advisors can alter investor behavior, suggesting that a more human-like presentation may shape reliance rather than simply improve understanding [20].

This is an important tension for financial advice because a system that feels more social and responsive may encourage stronger compliance even when the investor has not fully evaluated the recommendation, while comparable evidence from Italy further suggests that financial literacy and the continuing demand for human financial advice shape the use of robo-advisory services, reinforcing the importance of context in adoption research [27]. The value of conversational AI therefore lies not only in higher engagement, but in whether the added social presence improves informed understanding instead of passive acceptance.

The literature therefore points to a common design problem rather than separate adoption factors: in AI-enabled robo-advisory, deeper personalization may improve perceived suitability and trust-related evaluations while increasing the salience of data use and privacy concern, and conversational AI may improve explanation and social presence while strengthening reliance on automated recommendations. The unresolved issue is how investors respond when the data architecture of advice and the interaction interface are varied together while the recommendation itself remains constant.

2.5. Conceptual Model, Hypotheses, and Exploratory Questions

Prior studies show that trust, usefulness, personalization, expertise, privacy, and preferences for human or hybrid advice influence robo-advisor acceptance; however, less is known about how investors respond when the architecture of automated advice itself is experimentally varied. This study addresses that design-based gap by examining two features that change different parts of the advisory process: financial digital twins, which alter how the investor is represented, and conversational AI, which alters how the recommendation is explained. Because the same portfolio recommendation is held constant across all conditions, the analysis focuses on the perceived effects of advisory design rather than differences in the financial content of the advice.

Figure 1 presents the conceptual model guiding the analysis.

The model separates the most direct and theoretically grounded effects of the experimental treatments from downstream exploratory relationships. Digital-twin personalization is expected to increase perceived personalization by presenting the advisory system as using a richer and more adaptive representation of the investor's financial life, while also increasing privacy concern because deeper personalization implies broader use of personal and financial data. Conversational AI is expected to increase perceived social presence by shifting the advisory interface from a static dashboard to a more interactive explanation of the recommendation.




Figure 1. Conceptual model of responses to financial digital-twin and conversational-AI robo-advisory designs.

The exploratory part of the model covers downstream responses, including usefulness, trust-related evaluations, adoption intention, advice compliance, portfolio acceptance, and hypothetical allocation willingness, which remain relevant because robo-advisory adoption depends on perceived value, confidence, and willingness to rely on automated advice. Given the measurement diagnostics reported below, these constructs are not treated as validated latent variables, but as exploratory indicators of scenario-based investor responses. They are therefore used to describe suggestive patterns rather than to support confirmatory mechanism claims.

The confirmatory hypotheses are therefore limited to the relationships that are most directly tied to the experimental manipulations and supported by the strongest or most defensible measurement properties:

H1. *Financial digital-twin robo-advisors will produce higher perceived personalization than standard robo-advisors.*

H2. *Financial digital-twin robo-advisors will produce higher privacy concern than standard robo-advisors.*

H3. *Conversational AI interfaces will produce higher perceived social presence than plain dashboard interfaces.*

In addition to these hypotheses, the study examines three exploratory research questions that connect the experimental design to broader debates on trust, adoption, and behavioral reliance in automated financial advice:

RQ1. *How are digital-twin personalization and conversational AI associated with exploratory trust-related, usefulness-related, and willingness-related responses?*

RQ2. *Does the combined digital-twin and conversational-AI condition produce stronger exploratory indicators of behavioral willingness, including portfolio acceptance and hypothetical allocation willingness?*

RQ3. *Are the exploratory pathway patterns consistent with the proposed roles of perceived personalization, privacy concern, social presence, and trust-related responses in shaping stated willingness to rely on robo-advisory advice?*

The strongest inferential emphasis is placed on perceived personalization and privacy concern, while social presence is interpreted cautiously and downstream willingness-

related outcomes are treated as exploratory. The revised model therefore converts a broad theoretical debate into a more disciplined experimental test of the personalization–privacy trade-off in AI-enabled robo-advisory design, while preserving exploratory evidence on trust, adoption, and behavioral reliance for future research.

3. Materials and Methods

3.1. Research Design and Experimental Stimuli

This study used a randomized 2×2 between-subjects online experiment administered through the Qualtrics online survey platform (Qualtrics, Provo, UT, USA; web-based platform, version not applicable) to test two robo-advisory design features: personalization architecture and interface style. Personalization architecture varied between a standard robo-advisor profile and a financial digital-twin profile, while interface style varied between a plain dashboard and a conversational AI interface. Each participant evaluated only one fictional robo-advisory scenario, which avoided carryover effects and allowed responses to be compared across independently assigned experimental conditions.

The recommended portfolio was held constant across all four conditions at 60% global equity ETFs, 35% bond ETFs, and 5% cash equivalents, ensuring that observed differences reflected the advisory design rather than the financial content of the recommendation. The standard-profile scenarios described a robo-advisor using conventional onboarding information, including age, investment horizon, risk tolerance, income range, and investment goals, whereas the digital-twin scenarios described a dynamic, permission-based profile incorporating goals, income patterns, spending behavior, liabilities, household needs, account structure, tax context, and expected life events. The plain-dashboard scenarios presented the recommendation through a static interface, while the conversational-AI scenarios presented it through a chatbot-style explanation.

Table 1 summarizes the design.

Table 1. Experimental design, participant flow, and valid analytical sample.

Panel	Category	Digital Twin	Conversational AI	Description	Valid n
A. Participant flow	Initial survey responses	—	—	Responses obtained before screening and quality exclusions	411
A. Participant flow	Excluded responses	—	—	Screen-outs and exclusions for recall failure, incompleteness, or excessive missingness	75
A. Participant flow	Final analytical sample	—	—	Valid cases retained for analysis	336
B. Experimental cells	C1	No	No	Standard robo-advisor + plain dashboard	82
B. Experimental cells	C2	No	Yes	Standard robo-advisor + conversational AI chatbot	85
B. Experimental cells	C3	Yes	No	Digital-twin robo-advisor + plain dashboard	83
B. Experimental cells	C4	Yes	Yes	Digital-twin robo-advisor + conversational AI chatbot	86
B. Experimental cells	Total	—	—	All valid experimental conditions	336

Note: Excluded responses include screening failures and quality-control exclusions.

3.2. Recruitment, Eligibility, and Sample Characteristics

Participants were recruited online through the author’s professional and academic networks, producing a non-probability convenience sample of digitally reachable adults with varying investment experience and familiarity with digital financial services. Eligible respondents were required to be at least 18 years old and to report that they had personally invested or made investment decisions during the previous 24 months, a criterion intended to make the hypothetical robo-advisory evaluation more plausible.

A total of 411 responses were obtained; after screen-outs and exclusions for recall failure, incompleteness, or excessive missingness, the final analytical sample consisted of 336 valid respondents. Because recruitment relied on professional and academic networks,

investment experience was self-reported, and the sample was concentrated in Italy while including respondents from other countries, the findings are interpreted as evidence from a digitally reachable investor-experienced convenience sample rather than as representative evidence for the broader retail-investor population.

Baseline balance was assessed across demographic, investment-related, and digital-finance familiarity variables to evaluate whether random assignment produced comparable treatment groups. Because the diagnostics indicated differences in gender, investment experience, and baseline privacy concern, the unadjusted treatment-effect estimates were supplemented with covariate-adjusted models including these variables together with age, education, assets bracket, financial knowledge, and robo-advisor familiarity. These adjusted models were used to assess robustness at the observed baseline imbalance, not to replace the randomized experimental comparison. The randomized assignment provides the primary basis for causal comparison across experimental conditions, while covariate-adjusted models are used to assess robustness to observed baseline imbalance rather than to construct an observational identification strategy.

3.3. Procedure and Variable Measurement

After eligibility and consent were confirmed, participants completed baseline questions on investment experience, robo-advisor familiarity, prior robo-advisor use, chatbot familiarity, financial knowledge, and baseline privacy concern before being randomly assigned to one of the four experimental scenarios. They then completed manipulation-check items, answered post-exposure perceptual measures, made a hypothetical USD 10,000 allocation decision, responded to a portfolio-action item, and provided demographic information.

All perceptual and willingness-to-share-data items used seven-point response scales, with higher values indicating stronger agreement or greater willingness to allow the specified data source. Multi-item scores were calculated as arithmetic means, but their interpretation depended on the reliability diagnostics reported below; allocation willingness was measured as the dollar amount, from USD 0 to USD 10,000, that respondents would allow the robo-advisor to manage, while portfolio acceptance was coded as 1 for “accept as shown” or “accept with small changes” and 0 otherwise.

The measurement strategy distinguishes between confirmatory outcomes, cautiously interpreted outcomes, and exploratory indicators: perceived personalization and privacy concern are treated as the main interpretable composite outcomes, social presence is retained as a theoretically important but cautiously interpreted outcome, and trust, usefulness, adoption intention, advice compliance, fiduciary quality, data control/transparency, and data-sharing willingness are retained only as exploratory descriptive indicators because of their low internal consistency in this sample.

Table 2 summarizes the measurement hierarchy.

Table 2. Measurement hierarchy, coding, reliability, and analytical interpretation.

Construct/Variable	Items	Scale or Coding	Reliability/Status	Role in Revised Analysis
Digital-twin manipulation check	2	1–7 mean score	Manipulation diagnostic	Confirms recognition of digital-twin treatment
Conversational-AI manipulation check	2	1–7 mean score	Manipulation diagnostic	Confirms recognition of conversational-AI treatment
Perceived personalization	4	1–7 mean score	$\alpha = 0.77$	Main confirmatory outcome for H1
Privacy concern	4	1–7 mean score	$\alpha = 0.70$	Main confirmatory outcome for H2
Social presence	3	1–7 mean score	$\alpha = 0.69$	Cautiously interpreted outcome for H3

Table 2. *Cont.*

Construct/Variable	Items	Scale or Coding	Reliability/Status	Role in Revised Analysis
Usefulness/benefits	4	1–7 mean score	$\alpha = 0.50$	Exploratory response indicator
Trust	4	1–7 mean score	$\alpha = 0.43$	Exploratory trust-related indicator
Data control/transparency	3	1–7 mean score	$\alpha = 0.13$	Exploratory safeguard-perception indicator
Fiduciary quality	6	1–7 mean score	$\alpha = 0.34$	Exploratory trust-related design indicator
Adoption intention	3	1–7 mean score	$\alpha = 0.26$	Exploratory stated-willingness indicator
Advice compliance	3	1–7 mean score	$\alpha = 0.31$	Exploratory behavioral-reliance indicator
Data-sharing willingness	7	1–7 mean score	$\alpha = 0.40$	Exploratory privacy-related indicator
Allocation willingness	1	USD 0–10,000	Single-item bounded outcome	Exploratory hypothetical behavioral outcome
Portfolio acceptance	1	1 = accepted as shown or with small changes; 0 = otherwise	Binary outcome	Exploratory portfolio-action indicator

Note: Cronbach's alpha values are used as diagnostics for the interpretation of composite indicators. Low-reliability composites are not treated as validated latent constructs and are not used to support confirmatory mechanism claims.

The full questionnaire items are provided in Appendix A, Tables A1–A3, and the experimental scenario descriptions are provided in Appendix B, Table A4.

3.4. Measurement Reliability and Interpretation of Composite Indicators

Internal consistency diagnostics, reported in Table 2 and Appendix C, Table A7, were conducted before interpreting treatment effects, and the results established the hierarchy used throughout the analysis. Perceived personalization showed acceptable reliability, $\alpha = 0.77$, and privacy concern reached the conventional threshold, $\alpha = 0.70$, so these two constructs receive the strongest inferential emphasis; social presence was close to the threshold, $\alpha = 0.69$, and is interpreted cautiously because it remains theoretically tied to the conversational-AI manipulation but falls slightly below the standard benchmark.

Other composite indicators displayed low internal consistency, including trust, $\alpha = 0.43$; data control/transparency, $\alpha = 0.13$; usefulness, $\alpha = 0.50$; fiduciary quality, $\alpha = 0.34$; adoption intention, $\alpha = 0.26$; advice compliance, $\alpha = 0.31$; and data-sharing willingness, $\alpha = 0.40$. These measures are therefore not treated as validated latent constructs or as evidence of confirmed psychological mechanisms but are retained because they remain substantively relevant to robo-advisory acceptance and help preserve transparency about the original measurement design.

Accordingly, results involving these indicators, including regression coefficients and bootstrapped indirect effects, are interpreted only as exploratory treatment-related response patterns. They are used to describe suggestive scenario-based associations and to motivate future research with stronger measurement instruments, not to support confirmatory construct, mediation, or mechanism claims.

Because several downstream composites did not meet acceptable reliability thresholds, the revised analysis does not use CFA or SEM as confirmatory evidence, since doing so would imply a validated latent-variable structure that the measurement diagnostics do not support.

3.5. Analytical Strategy

The empirical analysis combined descriptive statistics, baseline balance checks, manipulation checks, reliability diagnostics, factorial treatment-effect models, robustness checks, and exploratory pathway tests. Balance tests used ANOVA for ordinal or approximately continuous variables and chi-square tests for nominal variables, while manipulation checks assessed whether participants recognized the digital-twin and conversational-AI features

as intended. Reliability diagnostics established the interpretation hierarchy for composite measures, treatment effects were estimated using models aligned with the 2×2 design, and robustness checks examined stability after covariate adjustment and across selected subsamples; exploratory pathway tests were reported only as suggestive evidence because several downstream composites displayed weak internal consistency. All statistical analyses were conducted using IBM SPSS Statistics for Windows, Version 29.0 (IBM Corp., Armonk, NY, USA).

The primary treatment-effect model was specified as follows:

$$Y_i = \beta_0 + \beta_1 DT_i + \beta_2 CAI_i + \beta_3 (DT_i \times CAI_i) + \gamma X_i + \varepsilon_i \quad (1)$$

In this model, Y_i denotes the outcome for respondent i , DT_i indicates assignment to a digital-twin condition, CAI_i indicates assignment to a conversational-AI condition, $DT_i \times CAI_i$ captures the interaction between the two experimental factors, X_i represents the baseline covariate vector used in adjusted models, and ε_i is the error term. The unadjusted specification estimates the randomized factorial contrasts, while the adjusted specification adds age, gender, education, assets bracket, investment experience, financial knowledge, baseline privacy concern, and robo-advisor familiarity to assess robustness to observed baseline imbalance and improve precision, particularly for gender, investment experience, and baseline privacy concern.

Treatment effects on Likert-type composite outcomes were estimated using two-way ANOVA and OLS regression with robust standard errors, while binary portfolio acceptance was analyzed through logistic regression with model fit assessed by the likelihood-ratio chi-square statistic. Allocation willingness, measured from USD 0 to USD 10,000, was analyzed as a bounded continuous outcome using OLS with robust standard errors and a nonparametric Kruskal–Wallis test, while Pearson correlations were used for approximately continuous composite variables and Spearman correlations were estimated as a robustness check for ordinal or bounded measures. Exploratory indirect-effect and pathway analyses used 5000 bootstrap resamples, but these results are interpreted only as preliminary pathway evidence because the relevant downstream composites, especially trust, adoption intention, and advice compliance, did not display adequate internal consistency.

Table 3 summarizes the analytical strategy.

Table 3. Analytical strategy by variable type.

Analytical Issue	Variables/Outcomes	Main Method	Robustness or Diagnostic Check	Interpretation
Randomized treatment effects	Main and exploratory composites	Two-way ANOVA; OLS with robust standard errors	Covariate-adjusted factorial model	Main evidence for H1–H2; cautious evidence for H3; exploratory evidence for weak composites
Baseline imbalance	Baseline covariates	Balance tests using ANOVA or chi-square	Covariate-adjusted models including baseline covariates	Assesses whether findings remain stable after observed baseline imbalance
Manipulation validity	Digital-twin recognition; conversational-AI recognition	Mean comparisons; treatment-effect regressions	Condition-level descriptive checks	Confirms that participants recognized the intended experimental treatments

Table 3. *Cont.*

Analytical Issue	Variables/Outcomes	Main Method	Robustness or Diagnostic Check	Interpretation
Measurement reliability	Multi-item composite indicators	Cronbach's alpha diagnostics	Interpretation hierarchy; exploratory-only reporting for weak composites	Prevents low-reliability composites from being treated as validated latent constructs
Binary outcome	Portfolio acceptance	Logistic regression	Likelihood-ratio chi-square; odds ratios	Exploratory portfolio-action indicator
Bounded monetary outcome	Allocation willingness, USD 0–10,000	OLS with robust standard errors	Kruskal–Wallis test across the four experimental conditions	Exploratory hypothetical allocation outcome
Correlation method	Composite, ordinal, and bounded variables	Pearson correlations for approx. continuous composite variables	Spearman correlations for ordinal or bounded robustness checks	Clarifies correlation method and scale treatment
Exploratory pathway tests	Exploratory pathway variables	Bootstrapped indirect effects with 5000 resamples	Exploratory-only interpretation	Suggestive pathway evidence, not confirmed mechanism

Note: Robust standard errors are reported where applicable. H1–H2 are treated as confirmatory; H3 is interpreted cautiously because social presence was borderline in reliability; and exploratory composites and indirect-effect pathways are reported only as suggestive evidence.

3.6. Ethical Considerations and Data Availability

Participation was voluntary and anonymous, and respondents were informed that their answers would be used for academic research and that they could withdraw before submitting the survey. No personally identifying information was collected, the study relied on hypothetical robo-advisory scenarios rather than actual investment decisions or real financial-account access, and formal ethical approval was not required under the applicable institutional requirements for anonymous adult survey research without personally identifying financial information.

The anonymized dataset supporting the study has been deposited in Mendeley Data under the title “Digital Twin Robo-Advisor Adoption Dataset V_2” and is publicly available at <https://doi.org/10.17632/39rw5ywj8r.1> [28]. The survey instrument, experimental vignettes, codebook, construct diagnostics, and robustness checks are provided in Appendices A–D.

4. Results

4.1. Sample and Baseline Balance

The final analyzed sample consisted of 336 respondents distributed relatively evenly across the four experimental conditions. There were 82 respondents in the standard robo-advisor and plain-dashboard condition, 85 in the standard robo-advisor and chatbot condition, 83 in the digital-twin and plain-dashboard condition, and 86 in the digital-twin and chatbot condition. Participants were adults with self-reported investment experience, recruited through the author's professional and academic networks. The sample was concentrated in Italy, which accounted for 202 respondents, or 60.1% of the sample, while the remaining respondents came from the United States, India, other European countries,

and selected Asian countries. Most participants were between 25 and 44 years old, and most reported at least some investment experience.

Table 4 reports the sample profile and baseline balance diagnostics.

Table 4. Sample profile and baseline balance.

Variable	Full Sample Statistic	<i>p</i> -Value
Valid analyzed sample	N = 336	n.a.
Condition sizes	C1 = 82; C2 = 85; C3 = 83; C4 = 86	n.a.
Age group	25–34: 100 (29.8%); 35–44: 98 (29.2%)	0.271
Gender	Female: 166 (49.4%); male: 155 (46.1%)	0.011
Education	Bachelor’s degree or above: 237 (70.5%)	0.651
Employment	Full-time: 173 (51.5%)	0.437
Country/region	Italy: 202 (60.1%); other countries: 134 (39.9%)	0.849
Assets bracket	Mean bracket: 2.82 (1.30)	0.914
Investment experience	2.73 (1.09)	0.002
Robo-advisor familiarity	3.90 (1.30)	0.221
Chatbot familiarity	4.99 (1.15)	0.193
Financial knowledge	3.99 (1.06)	0.694
Baseline privacy concern	4.27 (1.23)	0.031
Prior robo-advisor use	No prior use: 189 (56.2%)	0.406

Note: C1 = standard robo-advisor with plain dashboard; C2 = standard robo-advisor with chatbot interface; C3 = digital-twin robo-advisor with plain dashboard; C4 = digital-twin robo-advisor with chatbot interface. Balance *p*-values are based on ANOVA for ordinal or scale variables and chi-square tests for nominal variables. For nominal variables, the descriptive entry reports a compact summary, while the *p*-value is calculated using the full category distribution. Accordingly, the gender balance test includes all reported gender categories.

The balance checks indicate that the treatment groups were broadly similar in age, education, employment status, country or region, assets bracket, robo-advisor familiarity, chatbot familiarity, financial knowledge, and prior robo-advisor use. Statistically significant differences were observed for gender, investment experience, and baseline privacy concern, which are relevant because these variables may plausibly influence responses to automated financial advice, privacy-sensitive personalization, and AI-mediated interaction. The main treatment-effect estimates are therefore supplemented by covariate-adjusted models that include these variables, together with the other baseline controls, to assess whether the findings remain stable after accounting for observed baseline differences across conditions.

Consistent with the measurement hierarchy defined in Section 3.4, perceived personalization and privacy concern receive the strongest inferential emphasis because they display acceptable internal consistency, while social presence is interpreted cautiously because

its reliability was close to the conventional threshold. Trust, data control/transparency, usefulness, fiduciary quality, adoption intention, advice compliance, and data-sharing willingness are treated as exploratory descriptive indicators rather than validated latent constructs. This distinction reflects the purpose of the study, which is to provide evidence on scenario-based investor responses to AI-enabled robo-advisory design rather than to validate a full psychometric model of robo-advisor adoption.

4.2. Manipulation Checks and Descriptive Outcomes

The manipulation checks indicate that respondents recognized the intended experimental features. Digital-twin recognition was higher in the two digital-twin conditions than in the two standard robo-advisor conditions, with an average digital-twin effect of 2.30 points, $p < 0.001$. Chatbot recognition was higher in the two conversational-AI conditions than in the two plain-dashboard conditions, with an average conversational-AI effect of 2.08 points, $p < 0.001$.

The descriptive outcome pattern is reported in Table 5.

Table 5. Manipulation checks and outcome means by condition.

Outcome	C1 Standard + Plain	C2 Standard + Chatbot	C3 Digital Twin + Plain	C4 Digital Twin + Chatbot
Digital-twin recognition	2.82	3.03	5.09	5.35
Chatbot recognition	3.23	5.36	3.34	5.37
Perceived personalization	3.72	3.96	5.15	5.40
Social presence	3.38	4.84	3.65	5.01
Trust	4.12	4.41	4.40	4.78
Privacy concern	3.91	4.03	4.83	4.74
Usefulness	4.13	4.35	4.80	4.99
Adoption intention	4.08	4.32	4.31	4.66
Advice compliance	4.00	4.16	4.11	4.61
Allocation willingness, USD	3612	4304	4735	5933
Portfolio acceptance rate	43.9%	52.9%	60.2%	72.1%

Note: Scale outcomes are condition means on 1–7 scales. Portfolio acceptance includes respondents who accepted the recommendation as shown or with small changes. Because some composite outcomes displayed limited internal consistency, descriptive means for trust, usefulness, adoption intention, advice compliance, and related constructs should be interpreted as exploratory.

Consistent with the experimental manipulations, the digital-twin conditions showed higher mean levels of perceived personalization and privacy concern, the two composite outcomes with acceptable internal consistency. Usefulness also followed a positive descriptive pattern in the digital-twin conditions, although this result is interpreted cautiously because the usefulness index displayed limited internal consistency. The conversational-AI conditions showed higher social presence, which is also interpreted cautiously because its reliability estimate was close to, but slightly below, the conventional threshold. Secondary design-perception and privacy-related outcomes, including fiduciary quality, data control/transparency, and data-sharing willingness, are reported in Appendix C, Table A8.

4.3. Treatment Effects

Table 6 demonstrates the two-way treatment-effect regressions.

These models estimate randomized factorial treatment contrasts with robust standard errors and are interpreted according to the reliability hierarchy defined in Section 3.4.

Digital-twin personalization had a strong positive effect on perceived personalization, $b = 1.43$, and also increased privacy concern. These two findings carry the strongest inferential weight because both constructs displayed acceptable internal consistency. Usefulness, trust, adoption intention, and advice compliance also showed positive treatment

patterns, but these composite-based results should be interpreted as exploratory where internal consistency was weak. Hypothetical allocation willingness also increased under the digital-twin condition. This pattern supports H1 and H2, while the positive effects on usefulness, trust, adoption intention, advice compliance, and allocation willingness are interpreted as exploratory patterns. More broadly, the results suggest that the perceived benefits of richer personalization appear alongside higher privacy concern.

Table 6. Robust two-way treatment-effect regressions.

Predictor	Perceived Personalization	Social Presence	Trust	Privacy Concern	Usefulness	Adoption Intention	Advice Compliance	Allocation Willingness
DT effect	1.43 *** (0.05)	0.22 *** (0.05)	0.33 *** (0.05)	0.81 *** (0.07)	0.66 *** (0.05)	0.28 *** (0.06)	0.28 *** (0.06)	1376 *** (114)
CAI effect	0.24 *** (0.05)	1.41 *** (0.05)	0.33 *** (0.05)	0.01 (0.07)	0.20 *** (0.05)	0.29 *** (0.06)	0.33 *** (0.06)	945 *** (114)
DT x CAI	0.01 (0.10)	−0.09 (0.11)	0.09 (0.10)	−0.22 (0.14)	−0.03 (0.10)	0.11 (0.12)	0.34 ** (0.12)	506 * (228)
R ²	0.725	0.683	0.206	0.284	0.375	0.123	0.155	0.400
N	336	336	336	336	336	336	336	336

Note: Coefficients are reported with robust standard errors in parentheses. DT = digital twin; CAI = conversational AI. Treatment indicators were mean-centered, so DT and CAI coefficients represent average factorial effects across the other treatment factor. All outcomes except allocation willingness are measured on 1–7 scales. H1–H2 are confirmatory; H3 is interpreted cautiously; and coefficients for trust, usefulness, adoption intention, advice compliance, and related downstream composites are treated as exploratory because of low internal consistency. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Conversational AI had its largest effect on social presence, $b = 1.41$, providing cautious support for H3 because the social-presence index was close to, but slightly below, the conventional reliability threshold. The interaction term was positive and significant for advice compliance and allocation willingness, providing exploratory evidence relevant to RQ2 and suggesting that the combined digital-twin and conversational-AI condition was associated with stronger selected behavioral-willingness responses, especially allocation willingness and the exploratory advice-compliance index.

4.4. Exploratory Pathway Tests and Sensitivity Checks

The pathway evidence is interpreted as exploratory because the trust, adoption-intention, and advice-compliance indices displayed weak internal consistency. In the supplemental model using the exploratory trust index, the perceived personalization, usefulness, and social presence were positively associated with trust-related responses, with coefficients of $b = 0.14$, $p < 0.001$; $b = 0.12$, $p = 0.019$; and $b = 0.20$, $p < 0.001$, respectively. The exploratory trust index was also positively associated with adoption intention, $b = 0.24$, $p < 0.001$, and advice compliance, $b = 0.15$, $p = 0.019$, when the exploratory pathway variables were estimated jointly. These associations are therefore reported as suggestive pathway evidence rather than as confirmation of a validated psychological mechanism.

The bootstrapped indirect effects in Table 7 provide exploratory evidence consistent with a positive digital-twin pathway.

The privacy pathway was negative, as expected, but its confidence interval crossed zero and is therefore interpreted as directional rather than confirmatory evidence. Appendix D, Table A9 reports the robustness checks, including the binary portfolio-acceptance model, for which the likelihood-ratio model diagnostic was significant, $\chi^2(3, N = 336) = 14.940$, $p = 0.002$. These checks address the baseline imbalances in gender, investment experience, and baseline privacy concern by estimating covariate-adjusted models that include the digital-twin condition, conversational-AI condition, their interaction,

and controls for gender, investment experience, baseline privacy concern, age, education, assets bracket, financial knowledge, and robo-advisor familiarity.

Table 7. Exploratory pathway results.

Model	Pathway	Indirect Effect	Bootstrap SE	95% CI	Interpretation
A	DT → personalization → trust → adoption intention	0.060	0.025	[0.017, 0.111]	Exploratory support
B	DT → privacy concern → adoption intention	−0.068	0.039	[−0.146, 0.006]	Directional, not confirmatory
C	CAI → social presence → trust → advice compliance	0.050	0.024	[0.008, 0.103]	Exploratory support

Note: Indirect effects were estimated using 5000 bootstrap resamples. Because the trust, adoption-intention, and advice-compliance indices displayed weak internal consistency, indirect-effect results are interpreted as exploratory pathway evidence rather than as confirmation of validated psychological mechanisms.

The main effects for perceived personalization and privacy concern remain substantively similar after adjustment, suggesting that the central findings are not driven solely by the observed baseline imbalances. The broader pattern of findings also remains broadly stable for social presence and allocation willingness, while adoption-related, trust-related, and exploratory pathway findings continue to be interpreted cautiously because of the measurement limitations discussed above.

5. Discussion

The results clarify how two AI-enabled robo-advisory design features shape participants' responses when the underlying portfolio recommendation is held constant. Digital-twin personalization produced the clearest evidence by increasing both perceived personalization and privacy concern, the two composite outcomes with acceptable internal consistency, thereby supporting H1 and H2 and showing that participants recognized the added relevance of a richer investor profile while also perceiving its broader data requirements as more privacy-sensitive. Conversational AI mainly affected the perceived interactional quality of the advisory experience, with its strongest effect observed for social presence; because this index was close to, but slightly below, the conventional reliability threshold, the finding provides cautious support for H3.

The findings for trust, usefulness, adoption intention, advice compliance, and related composite indicators require a more restrained interpretation because, although they showed positive descriptive or regression-based patterns, their internal consistency was weak. These outcomes are therefore best treated as exploratory indicators of scenario-based responses rather than validated psychological constructs, meaning that the study provides experimental evidence on how investment-experienced participants evaluated described robo-advisory scenarios, not evidence of actual adoption behavior in a live advisory environment.

The combined digital-twin and conversational-AI condition produced the highest descriptive levels of several adoption-related and reliance-related indicators, including hypothetical allocation willingness and portfolio acceptance, while the regression interaction was clearest for allocation willingness and the exploratory advice-compliance index. This provides exploratory evidence relevant to RQ2 and suggests that richer personalization paired with conversational explanation may be especially relevant when participants evaluate whether to follow or allocate funds to a specific recommendation, although the evidence does not support a broad claim that the combined design uniformly increases adoption across all outcomes.

5.1. Scientific and Theoretical Implications

These results extend robo-advisor acceptance research by showing that responses to automated financial advice depend not only on general attitudes toward digital finance, but also on the architecture through which advice is generated and explained. While prior studies emphasize usefulness, trust, familiarity, provider credibility, and perceived expertise as drivers of robo-advisor acceptance [1–4,14–18], this study shows that perceived suitability also depends on how the investor is represented; a digital-twin profile may be evaluated more favorably than a standard profile because it captures a broader and more adaptive view of the investor's financial life, yet this benefit is accompanied by higher privacy concern.

Existing work has discussed financial digital twins as a promising direction for personalization, simulation, and decision support in financial services [5–7,24], but empirical evidence on user responses to this architecture remains limited. The present experiment suggests that participants can distinguish between a conventional onboarding profile and a more data-rich financial twin, and that this distinction affects perceived personalization and privacy concern. Although the evidence for usefulness, trust-related responses, and adoption-related willingness remains exploratory, it supports the plausible interpretation that financial digital twins may improve the perceived relevance of automated advice while increasing concern about the amount and sensitivity of personal data required to support that relevance.

Consistent with privacy-calculus logic [13], the personalization–privacy tension appears especially salient in robo-advisory, where platforms may request information about income, assets, liabilities, spending patterns, goals, household circumstances, tax position, and future life events. The contribution of the present study is to locate this logic within AI-enabled advisory architecture, showing that a richer investor representation can increase perceived suitability and relevance while also making the scope, sensitivity, and governance of data use more visible to users. This suggests that financial digital twins are likely to be more acceptable when users understand what data are collected, why those data are needed, how they influence the recommendation, and how access can be reviewed, corrected, limited, or withdrawn; without such safeguards, deeper personalization may improve perceived relevance while also making users feel overexposed.

The conversational-AI findings contribute to research on social presence and automated financial advice by showing that the chatbot condition increased perceived social presence and was associated with more favorable patterns in several stated willingness outcomes, consistent with evidence that conversational robo-advisors can make onboarding and recommendation explanation feel more interactive [8,22,23]. Because several downstream composite measures displayed weak internal consistency, these results should not be interpreted as evidence that conversational AI reliably increases trust or adoption in a general sense, but as more cautious evidence that it strengthens the interactional experience and may support selected willingness responses, especially when paired with a richer personalization architecture.

Prior research has warned that social design elements in robo-advisory can shape investor behavior and increase reliance on automated recommendations [20], a concern reflected in the present study by the stronger allocation willingness and positive exploratory advice-compliance pattern observed in the combined digital-twin and conversational-AI condition. This suggests that conversational explanation may support confidence in the recommendation, but also creates a design responsibility: in financial advice, the objective should not be to maximize comfort, fluency, or compliance for its own sake, but to provide conversational explainability that helps users understand why a recommendation

is suitable, what assumptions it depends on, what risks it involves, and what alternatives remain available.

5.2. Practical and Regulatory Implications

Because the evidence comes from scenario-based evaluations, these implications should be read as design guidance rather than as evidence of observed user behavior in functioning robo-advisory systems.

The practical implication is that robo-advisory providers should present financial digital twins not as expanded data-collection systems, but as controlled personalization tools that improve suitability, transparency, and user understanding. Users should be able to inspect the financial profile used by the system, identify the assumptions influencing the recommendation, update or correct profile information, and manage data permissions in a visible and reversible way, because the more comprehensive the digital twin becomes, the more important it is to make the profile auditable from the user's perspective.

Conversational AI should follow the same principle by clarifying the recommendation, explaining portfolio allocation, identifying relevant risks, and supporting comparison with alternatives, rather than merely increasing comfort or stated willingness in ways that may be problematic in financial advice if understanding does not improve. The better design objective is to make conversational AI an explanation layer rather than a persuasion layer, especially because the strongest willingness-related responses appeared when it was combined with digital-twin personalization.

The findings also have regulatory and consumer-protection implications because financial digital twins depend on broader and potentially more continuous data flows than standard robo-advisory questionnaires. Consent, data minimization, explainability, and user control should therefore apply not only to the algorithm that generates the recommendation, but also to the data architecture that feeds it and the interface that frames it. Conversational AI raises a related concern because a more socially present interface may increase stated willingness to follow automated advice, making it important to distinguish explanation that improves understanding from interaction design that encourages passive acceptance.

5.3. Limitations and Future Research

The findings should be interpreted in light of several limitations, beginning with the use of scenario-based descriptions rather than a live robo-advisory platform, which allowed clean manipulation of digital-twin personalization and conversational AI but could not reproduce all features of real platform use. The findings should therefore be understood as design-level experimental evidence about perceived advisory architectures, not as field evidence of revealed investor behavior. Allocation willingness was hypothetical, portfolio acceptance was based on stated response rather than actual investment behavior, and the sample was recruited through professional and academic networks, with self-reported investment experience and concentration in Italy despite including respondents from several countries.

Measurement reliability is a further limitation because several multi-item indices, including trust, usefulness, fiduciary quality, adoption intention, advice compliance, data control/transparency, and data-sharing willingness, displayed low internal consistency; accordingly, the revised interpretation places the greatest weight on perceived personalization and privacy concern, interprets social presence cautiously, and treats the remaining composite indicators as exploratory.

The results should therefore be interpreted as evidence from a digitally reachable investor-experienced convenience sample rather than as representative evidence for the

broader retail-investor population. Future studies should use more extensively validated scales, conduct confirmatory factor analysis where sample size permits, and supplement stated willingness measures with behavioral or incentivized outcomes.

Future research should test these design features in functioning robo-advisory environments where users can interact with dashboards, chatbots, account-linking tools, consent interfaces, and explanation modules. Field experiments with FinTech providers would be especially useful for observing real data-sharing choices, portfolio acceptance, allocation decisions, and continued use. Further studies could also compare different consent dashboards, privacy labels, explanation formats, and warnings against overreliance. Cross-country research would help determine whether responses to financial digital twins and conversational AI vary across regulatory regimes, financial-literacy levels, and cultural attitudes toward automated advice.

6. Conclusions

This study shows that robo-advisory design architecture shapes perceived personalization, privacy concern, and selected indicators of stated behavioral willingness in a scenario-based experiment. By holding the portfolio recommendation constant, digital-twin personalization produced the clearest effects by increasing both perceived personalization and privacy concern, indicating that a more adaptive investor profile can make automated advice appear more relevant while also making its data requirements more salient. Conversational AI strengthened perceived interactional quality mainly through social presence, although this result is interpreted cautiously because the index was close to, but slightly below, the conventional reliability threshold. Accordingly, the strongest conclusion concerns the coexistence of higher perceived personalization and higher privacy concern, while findings involving trust, usefulness, adoption intention, advice compliance, and allocation willingness remain exploratory.

The central conclusion is that financial digital twins and conversational AI may improve how participants evaluate robo-advisory services, but only when personalization, explanation, and data use are governed in ways that users can understand and control. A richer investor profile can increase perceived relevance, yet it also raises privacy concern, while a conversational interface can make advice easier to engage with but may also encourage stronger reliance on automated recommendations. Exploratory willingness-related patterns appeared most clearly when digital-twin personalization and conversational AI were combined, especially for allocation willingness and the exploratory advice-compliance index, but these findings should be interpreted as scenario-based responses rather than as evidence of actual adoption behavior in a live robo-advisory environment.

The study contributes to behavioral FinTech research by providing experimental evidence on two next-generation robo-advisory features that have often been discussed conceptually but tested less directly. Theoretically, it shows that AI-enabled robo-advisory design should be understood through the joint lens of personalization, privacy concern, and interaction quality rather than through adoption intention alone. Practically, it suggests that future robo-advisors should be designed not only for optimization and convenience, but also for explainability, consent, data control, and active investor understanding.

Funding: This research received no external funding.

Institutional Review Board Statement: Ethical review and approval were waived for this study under applicable institutional requirements for anonymous survey research involving adult participants, with no collection of personally identifying information.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The anonymized dataset supporting the findings of this study is publicly available in Mendeley Data under the title “Digital Twin Robo-Advisor Adoption Dataset V_2” at <https://doi.org/10.17632/39rw5ywj8r1>.

Acknowledgments: The author thanks the respondents who participated in the survey and the professional contacts who supported its dissemination. During the preparation of this manuscript, the author used generative AI support for language polishing, formatting, and structural refinement. The author reviewed and edited all outputs and takes full responsibility for the content of this publication.

Conflicts of Interest: The author declares no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
ANOVA	Analysis of Variance
CAI	Conversational Artificial Intelligence
CI	Confidence Interval
DT	Digital Twin
ETF	Exchange-Traded Fund
OLS	Ordinary Least Squares
SE	Standard Error
SPSS	Statistical Package for the Social Sciences
USD	United States dollar
UTAUT	Unified Theory of Acceptance and Use of Technology

Appendix A. Survey Questionnaire

Unless otherwise stated, all perceptual items used seven-point Likert scales, where 1 = strongly disagree and 7 = strongly agree.

Table A1. Eligibility and baseline items.

Block	Item	Response Format
Consent	Do you agree to participate in this study?	1 = yes; 0 = no
Eligibility	Are you at least 18 years old?	1 = yes; 0 = no
Eligibility	Have you personally invested or made investment decisions during the last 24 months?	1 = yes; 0 = no
Baseline	How familiar are you with robo-advisors or automated investment platforms?	1 = not at all familiar; 7 = very familiar
Baseline	Have you ever used a robo-advisor or automated investment platform?	1 = no; 2 = tried once; 3 = previously used; 4 = currently use
Baseline	How would you rate your investment experience?	1 = novice; 2 = beginner; 3 = intermediate; 4 = advanced; 5 = professional
Baseline	How familiar are you with AI chatbots or conversational AI tools?	1 = not at all familiar; 7 = very familiar
Baseline	How would you rate your overall financial knowledge?	1 = very low; 7 = very high
Baseline	In general, I am concerned about how financial apps use my personal data.	1 = strongly disagree; 7 = strongly agree

Table A2. Manipulation-check and post-exposure perceptual items.

Construct	Item
Digital-twin recognition	The advisor appeared to create a detailed digital profile of the investor.
Digital-twin recognition	The advisor seemed to use ongoing personal financial data to personalize the recommendation.
Conversational-AI recognition	The advisor interacted in a conversational way.
Conversational-AI recognition	The advisor felt more like a chat-based assistant than a static dashboard.
Perceived personalization	The recommendation felt tailored to the investor's specific financial situation.
Perceived personalization	The advisor appeared to consider future life changes, not only current risk tolerance.
Perceived personalization	The advisor seemed capable of updating advice as the investor's circumstances change.
Perceived personalization	The advisor seemed to understand the investor as a whole financial person rather than as a simple questionnaire profile.
Social presence	The advisor's communication felt interactive.
Social presence	The advisor made the recommendation easier to understand.
Social presence	The advisor felt responsive to possible investor questions.
Trust	I would trust this robo-advisor to provide reliable investment guidance.
Trust	This robo-advisor seems competent in making portfolio recommendations.
Trust	This robo-advisor seems honest about how it generates advice.
Trust	I would feel confident relying on this robo-advisor for basic investment guidance.
Privacy concern	I would be concerned about how much personal data this service uses.
Privacy concern	The personalization seems to require sensitive financial information.
Privacy concern	I would worry that my data could be misused or shared in ways I do not expect.
Privacy concern	The benefits of personalization may not justify the privacy risks.
Data control/transparency	The service seems transparent about what data it uses.
Data control/transparency	I would feel in control of the data used to generate recommendations.
Data control/transparency	The service seems to allow users to review or change the data used for personalization.
Usefulness	This robo-advisor would make investment planning more useful.
Usefulness	This robo-advisor would help investors make better long-term financial decisions.
Usefulness	This robo-advisor would save time compared with managing investments alone.
Usefulness	This robo-advisor offers clear value compared with a basic investment app.
Fiduciary quality	The advisor appears capable of managing portfolio risk responsibly.
Fiduciary quality	The advisor seems to consider the investor's financial well-being.
Fiduciary quality	The advisor appears to use enough information to make a suitable recommendation.
Fiduciary quality	The advisor seems aligned with the investor's interests.
Fiduciary quality	The advisor appears transparent about why the recommendation was made.
Fiduciary quality	The advisor does not seem designed mainly to benefit the platform at the investor's expense.

Table A3. Adoption, behavioral, data-sharing, and demographic items.

Block	Item	Response Format
Adoption intention	I would consider using this robo-advisor.	1 = strongly disagree; 7 = strongly agree
Adoption intention	I would be willing to open an account with a service like this.	1 = strongly disagree; 7 = strongly agree
Adoption intention	I would recommend a service like this to someone with investment needs.	1 = strongly disagree; 7 = strongly agree
Advice compliance	I would be comfortable following this portfolio recommendation.	1 = strongly disagree; 7 = strongly agree
Advice compliance	I would be likely to accept the recommendation without major changes.	1 = strongly disagree; 7 = strongly agree
Advice compliance	I would rely on this advisor's guidance during normal market conditions.	1 = strongly disagree; 7 = strongly agree
Recall check	According to the scenario, what portfolio was recommended?	1 = 60% global equity ETFs, 35% bond ETFs, 5% cash equivalents; 2 = 35% global equity ETFs, 60% bond ETFs, 5% cash equivalents; 3 = 100% cash equivalents; 4 = do not remember
Allocation willingness	Imagine you had USD 10,000 available to invest. How many dollars would you allow this robo-advisor to manage?	0 to 10,000
Portfolio action	What would you do with the recommended 60/35/5 portfolio?	1 = accept as shown; 2 = accept with small changes; 3 = make major changes; 4 = reject; 5 = need human advisor first
Data-sharing willingness	Bank transaction data	1 = definitely not allow; 7 = definitely allow
Data-sharing willingness	Brokerage or retirement account data	1 = definitely not allow; 7 = definitely allow
Data-sharing willingness	Income or payroll pattern data	1 = definitely not allow; 7 = definitely allow
Data-sharing willingness	Tax profile information	1 = definitely not allow; 7 = definitely allow
Data-sharing willingness	Spending category data	1 = definitely not allow; 7 = definitely allow
Data-sharing willingness	Household or family-planning information	1 = definitely not allow; 7 = definitely allow
Data-sharing willingness	Social-media signals, such as major life events	1 = definitely not allow; 7 = definitely allow
Demographics	What is your age group?	1 = 18–24; 2 = 25–34; 3 = 35–44; 4 = 45–54; 5 = 55–64; 6 = 65+
Demographics	What is your gender?	1 = female; 2 = male; 3 = non-binary; 4 = prefer not to say; 5 = self-describe
Demographics	What is your highest completed education?	1 = high school or less; 2 = some college; 3 = Bachelor's degree; 4 = Master's degree; 5 = doctoral/professional; 6 = prefer not to say
Demographics	What is your current employment status?	1 = full-time; 2 = part-time; 3 = self-employed; 4 = student; 5 = retired; 6 = unemployed; 7 = prefer not to say

Table A3. *Cont.*

Block	Item	Response Format
Demographics	Approximately how much do you have in investable assets, excluding your primary residence?	1 = less than USD 5000; 2 = USD 5000–24,999; 3 = USD 25,000–99,999; 4 = USD 100,000–249,999; 5 = USD 250,000 or more; 6 = prefer not to say
Demographics	What is your country or region of residence?	Open text, cleaned into country-region groups

Appendix B. Experimental Stimuli

All four conditions presented the same recommended portfolio: 60% global equity ETFs, 35% bond ETFs, and 5% cash equivalents. Only the personalization architecture and interface style varied across conditions.

Table A4. Experimental conditions and scenario descriptions.

Condition	Profile Description	Interface Description
C1 Standard + Plain	The robo-advisor used a standard questionnaire based on age, investment horizon, risk tolerance, income range, and investment goals.	The recommendation was shown through a plain dashboard with the portfolio allocation, a short risk description, and expected diversification benefits.
C2 Standard + Chatbot	The robo-advisor used a standard questionnaire based on age, investment horizon, risk tolerance income range, and investment goals.	The recommendation was explained through a conversational AI chatbot that described the role of equities, bonds, and cash equivalents in plain language.
C3 Digital Twin + Plain	The robo-advisor created a financial digital twin using a dynamic, permission-based profile that could include goals, income pattern, spending behavior, liabilities, household needs, account structure, tax context, and expected life events.	The recommendation was shown through a plain dashboard with the portfolio allocation, a short risk description, and expected diversification benefits.
C4 Digital Twin + Chatbot	The robo-advisor created a financial digital twin using a dynamic, permission-based profile that could include goals, income pattern, spending behavior, liabilities, household needs, account structure, tax context, and expected life events.	The recommendation was explained through a conversational AI chatbot that described how the digital twin profile informed the portfolio recommendation.

Appendix C. Codebook, Reliability Diagnostics, and Exploratory Outcomes

The main analytical variables, construct diagnostics, and selected secondary outcomes are summarized below. Table A5 reports the key analytical variables used in the experiment.

The reliability diagnostics are reported to clarify the interpretation hierarchy used in the analysis. Perceived personalization and privacy concern showed acceptable internal consistency and therefore received the strongest inferential weight. Social presence was close to the conventional reliability threshold and is interpreted cautiously. Trust, data control/transparency, usefulness, fiduciary quality, adoption intention, advice compliance, and data-sharing willingness displayed low alpha values and are therefore retained only as exploratory indicators of scenario-based responses rather than as validated latent constructs. The analysis consequently avoids relying on these weaker indices as primary evidence for

confirmatory mechanism claims and treats findings involving them as suggestive rather than confirmatory.

Table A5. Key analytical variables.

Variable	Description	Coding/Scale
Condition	Experimental condition	1 = C1; 2 = C2; 3 = C3; 4 = C4
DT	Digital-twin treatment	1 = digital twin; 0 = standard profile
CAI	Conversational-AI treatment	1 = chatbot; 0 = plain dashboard
Investment experience	Self-rated investment experience	1 = novice; 5 = professional
Robo-advisor familiarity	Familiarity with robo-advisors	1 = not at all familiar; 7 = very familiar
Chatbot familiarity	Familiarity with AI chatbots	1 = not at all familiar; 7 = very familiar
Financial knowledge	Self-rated financial knowledge	1 = very low; 7 = very high
Baseline privacy concern	General concern with financial app data use	1 = strongly disagree; 7 = strongly agree
Allocation willingness	Amount assigned to the robo-advisor	USD 0–10,000
Portfolio acceptance	Accepted as shown or with small changes	1 = accepted; 0 = otherwise
Digital-twin recognition	Mean of two manipulation-check items	1–7
Chatbot recognition	Mean of two manipulation-check items	1–7
Perceived personalization	Mean of four personalization items	1–7
Social presence	Mean of three social-presence items	1–7
Trust	Mean of four trust items	1–7
Privacy concern	Mean of four privacy-concern items	1–7
Data control/transparency	Mean of three data-control items	1–7
Usefulness	Mean of four usefulness items	1–7
Fiduciary quality	Mean of six fiduciary-quality items	1–7
Adoption intention	Mean of three adoption-intention items	1–7
Advice compliance	Mean of three advice-compliance items	1–7
Data-sharing willingness	Mean of seven data-sharing items	1–7

Table A6. Interpretation hierarchy based on reliability diagnostics.

Category	Construct	Interpretation
Main confirmatory composite outcomes	Perceived personalization; privacy concern	These constructs showed acceptable internal consistency and received the strongest inferential weight in the analysis.
Cautiously interpreted composite outcome	Social presence	This construct was close to the conventional reliability threshold and is interpreted cautiously.
Exploratory indicators	Trust; data control/transparency; usefulness; fiduciary quality; adoption intention; advice compliance; data-sharing willingness	Low internal consistency; exploratory indicators only, not validated latent constructs.

Table A7. Construct diagnostics.

Construct	Items	Mean	SD	Alpha
Perceived personalization	4	4.57	0.85	0.77
Social presence	3	4.23	0.87	0.69

Table A7. *Cont.*

Construct	Items	Mean	SD	Alpha
Trust	4	4.43	0.52	0.43
Privacy concern	4	4.38	0.77	0.70
Data control/transparency	3	4.25	0.50	0.13
Usefulness	4	4.57	0.56	0.50
Fiduciary quality	6	4.37	0.40	0.34
Adoption intention	3	4.34	0.59	0.26
Advice compliance	3	4.22	0.59	0.31
Data-sharing willingness	7	3.52	0.45	0.40

Note: Alpha values are reported as reliability diagnostics. Perceived personalization and privacy concern are treated as the main confirmatory composite outcomes, social presence is interpreted cautiously, and the remaining low-reliability composites are treated as exploratory indicators rather than validated latent constructs.

Table A8. Exploratory secondary outcomes.

Outcome	Full Sample Mean	SD	Interpretation
Data control/transparency	4.25	0.50	Moderate perceived control and transparency, interpreted as exploratory because of low internal consistency.
Fiduciary quality	4.37	0.40	Moderate perceived investor alignment, interpreted as exploratory because of low internal consistency.
Data-sharing willingness	3.52	0.45	Below-midpoint willingness to share additional data, interpreted as exploratory because of low internal consistency.

Appendix D. Robustness Checks

This appendix summarizes additional analyses used to assess the stability of the main findings, with particular attention to the observed baseline imbalances in gender, investment experience, and baseline privacy concern.

Table A9. Summary of robustness checks.

Check	Specification	Main result
Covariate-adjusted models addressing baseline imbalance	Controls for baseline imbalances and key respondent characteristics.	The adjusted results remained similar for perceived personalization and privacy concern, and broadly consistent for social presence and allocation willingness. Results for weaker composites remain exploratory.
No prior robo-advisor use	Subsample restricted to respondents with no prior use (N = 189).	DT and CAI effects remained positive for the main perceptual outcomes and selected exploratory willingness outcomes.
Italy-only subsample	Subsample restricted to Italy respondents (N = 202).	DT increased perceived personalization and privacy concern, CAI increased social presence, and both treatments remained positively associated with allocation willingness.
Allocation willingness	Kruskal–Wallis test across the four conditions.	$H = 139.26, p < 0.001$; the highest allocation willingness remained in C4.
Portfolio acceptance	Logistic regression on binary acceptance of the recommended portfolio.	Likelihood-ratio $\chi^2(3, N = 336) = 14.940, p = 0.002$; $OR_{DT} = 2.11, p = 0.001$; $OR_{CAI} = 1.57, p = 0.048$; $OR_{interaction} = 1.19, p = 0.706$.
Exploratory pathway tests	Bootstrapped indirect effects with 5000 resamples.	The pathway evidence remained exploratory, and the privacy pathway was negative but non-confirmatory.

References

1. Jung, D.; Dorner, V.; Glaser, F.; Morana, S. Robo-advisory: Digitalization and automation of financial advisory. *Bus. Inf. Syst. Eng.* **2018**, *60*, 81–86. [\[CrossRef\]](#)
2. Belanche, D.; Casaló, L.V.; Flavián, C. Artificial intelligence in FinTech: Understanding robo-advisors adoption among customers. *Ind. Manag. Data Syst.* **2019**, *119*, 1411–1430. [\[CrossRef\]](#)
3. D’Acunto, F.; Prabhala, N.; Rossi, A.G. The promises and pitfalls of robo-advising. *Rev. Financ. Stud.* **2019**, *32*, 1983–2020. [\[CrossRef\]](#)
4. Brenner, L.; Meyll, T. Robo-advisors: A substitute for human financial advice? *J. Behav. Exp. Financ.* **2020**, *25*, 100275. [\[CrossRef\]](#)
5. Anshari, M.; Almunawar, M.N.; Masri, M. Digital twin: Financial technology’s next frontier of robo-advisor. *J. Risk Financ. Manag.* **2022**, *15*, 163. [\[CrossRef\]](#)
6. Bonelli, M.I.; Döngül, E.S. Robo-advisors in the financial services industry: Recommendations for full-scale optimization, digital twin integration, and leveraging natural language processing trends. In *Proceedings of the 2023 9th International Conference on Virtual Reality (ICVR), Xianyang, China, 14 May 2023*; IEEE: Piscataway, NJ, USA, 2023; pp. 268–275. [\[CrossRef\]](#)
7. Bonelli, M.I.; Liu, J. Revolutionizing robo-advisors: Unveiling global financial markets, AI-driven innovations, and technological landscapes for enhanced investment decisions. *Adv. Sci. Technol. Eng. Syst. J.* **2024**, *9*, 33–44. [\[CrossRef\]](#)
8. Hildebrand, C.; Bergner, A. Conversational robo advisors as surrogates of trust: Onboarding experience, firm perception, and consumer financial decision making. *J. Acad. Mark. Sci.* **2021**, *49*, 659–676. [\[CrossRef\]](#)
9. Bol, N.; Dienlin, T.; Kruijkemeier, S.; Sax, M.; Boerman, S.C.; Strycharz, J.; Helberger, N.; de Vreese, C.H. Understanding the effects of personalization as a privacy calculus: Analyzing self-disclosure across health, news, and commerce contexts. *J. Comput.-Mediat. Commun.* **2018**, *23*, 370–388. [\[CrossRef\]](#)
10. Back, C.; Morana, S.; Spann, M. When do robo-advisors make us better investors? The impact of social design elements on investor behavior. *J. Behav. Exp. Econ.* **2023**, *103*, 101984. [\[CrossRef\]](#)
11. Davis, F.D. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *Manag. Inf. Syst. Q.* **1989**, *13*, 319–340. [\[CrossRef\]](#) [\[PubMed\]](#)
12. Komiak, S.Y.X.; Benbasat, I. The effects of personalization and familiarity on trust and adoption of recommendation agents. *Manag. Inf. Syst. Q.* **2006**, *30*, 941–960. [\[CrossRef\]](#)
13. Lourenço, C.J.S.; Dellaert, B.G.C.; Donkers, B. Whose algorithm says so: The relationships between type of firm, perceptions of trust and expertise, and the acceptance of financial robo-advice. *J. Interact. Mark.* **2020**, *49*, 107–124. [\[CrossRef\]](#)
14. Venkatesh, V.; Morris, M.G.; Davis, G.B.; Davis, F.D. User acceptance of information technology: Toward a unified view. *Manag. Inf. Syst. Q.* **2003**, *27*, 425–478. [\[CrossRef\]](#)
15. de Oliveira Santini, F.; Hasan Jafar, S.; Junior Ladeira, W.; Rasul, T. Drivers of robo-advisors adoption in financial services: A meta-analytical study. *Int. J. Bank Mark.* **2026**, *44*, 439–453. [\[CrossRef\]](#)
16. Mayer, R.C.; Davis, J.H.; Schoorman, F.D. An integrative model of organizational trust. *Acad. Manag. Rev.* **1995**, *20*, 709–734. [\[CrossRef\]](#)
17. Glikson, E.; Woolley, A.W. Human trust in artificial intelligence: Review of empirical research. *Acad. Manag. Ann.* **2020**, *14*, 627–660. [\[CrossRef\]](#)
18. Nourallah, M.; Öhman, P.; Amin, M. No trust, no use: How young retail investors build initial trust in financial robo-advisors. *J. Financ. Report. Account.* **2023**, *21*, 60–82. [\[CrossRef\]](#)
19. Cao, X.; de Zwaan, L.; Wong, V. Building trust in robo-advisory: Technology, firm-specific and system trust. *Qual. Res. Financ. Mark.* **2026**, *18*, 408–428. [\[CrossRef\]](#)
20. Luo, H.; Liu, X.; Lv, X.; Hu, Y.; Ahmad, A.J. Investors’ willingness to use robo-advisors: Extrapolating influencing factors based on the fiduciary duty of investment advisors. *Int. Rev. Econ. Financ.* **2024**, *94*, 103411. [\[CrossRef\]](#)
21. Fatima, S.; Chakraborty, M. Adoption of artificial intelligence in financial services: The case of robo-advisors in India. *IIMB Manag. Rev.* **2024**, *36*, 113–125. [\[CrossRef\]](#)
22. Bhat, A.K.; Krishnan, G. A synergistic review: Revolutionizing financial systems through digital twins. In *Proceedings of the 2025 3rd International Conference on Intelligent Cyber Physical Systems and Internet of Things (ICoICI), Coimbatore, India, 17–19 September 2025*; IEEE: Piscataway, NJ, USA, 2025; pp. 939–944. [\[CrossRef\]](#)
23. Tahvildari, M. Analysis of investor profiling and portfolio recommendations by robo-advisors in Germany. *Glob. Bus. Financ. Rev.* **2025**, *30*, 19–33. [\[CrossRef\]](#)
24. Tang, Y.; Son, H. How the sociality of AI digital human advisors shapes user experience value in digital finance: The mediating role of social presence. *J. Theor. Appl. Electron. Commer. Res.* **2026**, *21*, 79. [\[CrossRef\]](#)
25. Wang, C.; Aw, E.C. Trust (or not) the bot? Understanding recommendation adoption and stickiness to the robo-advisor. *Int. J. Bank Mark.* **2026**, in press. [\[CrossRef\]](#)
26. Brüggén, L.; Gianni, R.; de Haan, F.; Hogreve, J.; Meacham, D.; Post, T.; van der Werf, M. AI-based financial advice: An ethical discourse on AI-based financial advice and ethical reflection framework. *J. Public Policy Mark.* **2025**, *44*, 436–456. [\[CrossRef\]](#)

27. Aristei, D.; Gallo, M. Financial literacy, robo-advising, and the demand for human financial advice: Evidence from Italy. *J. Behav. Exp. Financ.* **2026**, *49*, 101125. [[CrossRef](#)]
28. Bonelli, M.I. *Digital Twin Robo-Advisor Adoption Dataset V_2*; V1; Mendeley Data: London, UK, 2026. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.