Journal Publication of International Research for Engineering and Management (JOIREM)
Volume: 04 Issue: 09 | Sep-2026
ISSN (O) 3107-6696

When AI Recommends, Who Decides? Exploring User Control and Decision-Making in AI-Generated Financial Recommendations

Anamika Yadav, Ruchi Gaur

M.Des Student, World University of Design, Rajiv Gandhi Education City, Sonipat, Haryana, India

anamikayadav.25@wud.ac.in

Professor, World University of Design, Rajiv Gandhi Education City, Sonipat, Haryana, India
ruchi@wud.ac.in

Abstract - This study examines how AI-generated financial recommendations and digital interface design influence users' trust, perceived control, decision confidence, and willingness to follow financial guidance. The study begins from a simple but important question: when an AI system recommends what a person should do with their money, who still feels responsible for the final decision? AI can process large amounts of information quickly and can present users with recommendations that appear clear and useful. However, the ease of following a recommendation does not necessarily mean that the recommendation has been understood, evaluated, or consciously accepted by the user. This makes the interface through which an AI recommendation is communicated an important part of the decision-making experience.

The proposed study uses a user-centered experimental approach with prototype digital interfaces that present the same type of AI-generated financial recommendation in different ways. Participants will complete simulated financial decision-making tasks and will be asked to evaluate the recommendation and make a final choice. Their experience will be examined through measures related to trust, perceived control, decision confidence, clarity, usability, and willingness to act on the recommendation. The study is designed to compare how different presentation approaches affect the movement from receiving AI guidance to making an independent decision.

The study does not assume that the purpose of an AI interface should be to increase acceptance of its recommendation. Instead, it considers appropriate decision support as an interaction in which the user can understand the recommendation, question it when necessary, compare it with alternatives, and retain a meaningful sense of control. Existing research has separately examined trust in AI, over-reliance, recommendation acceptance, autonomy, explainability, and financial AI adoption. This research brings these concerns together from an interface and user-experience perspective. The expected contribution is a clearer understanding of how interface choices can support AI as a helpful guide without making the user feel that the decision has already been made for them.

***Keywords*:** AI-assisted decision-making, financial recommendation systems, recommendation acceptance, user agency, perceived control, decision confidence, AI trust, interface psychology, human-AI interaction, financial UX.

Introduction

Making decisions about money is a common part of everyday life. People decide how much to save, how much to spend, whether a purchase is affordable, whether a financial option fits their situation, and whether a recommendation is worth following. Although these decisions may look routine, they often involve uncertainty, incomplete information, personal priorities, and the possibility of consequences that are not immediately visible. A person may therefore need more than information alone. They may need help understanding the situation and comparing possible choices while still feeling that the final decision belongs to them.

The growing use of artificial intelligence changes this experience. AI systems can interpret large amounts of information, identify patterns, summarize complex material, and generate recommendations. Instead of asking a person to examine every available option independently, the system can present a suggested direction. As Amershi et al. (2019) emphasize in their human-AI interaction guidelines, the way an AI system communicates and supports user action is part of the interaction itself. This can reduce effort and make decision-making feel faster. At the same time, it changes the role of the interface. The system is no longer simply showing information; it is influencing what the user notices, considers, and may eventually choose.

This shift raises the central question of this research: when AI tells people what they should do with their money, how much should they rely on it? The question is not intended to suggest that AI recommendations are automatically good or bad. Rather, it focuses attention on the interaction between the recommendation, the interface, and the person who must make the final decision. An AI recommendation can be technically sophisticated while still being difficult for a user to understand. Conversely, an interface can make a recommendation appear extremely convincing without giving the user enough opportunity to question it. This concern is consistent with research on explainability and human-AI

interaction, where understanding depends not only on whether an explanation exists but also on how it is communicated.

Research on human-AI decision-making has shown that people can sometimes over-rely on AI advice, including situations in which the recommendation conflicts with contextual information or with their own assessment. Klingbeil, Grützner, and Schreck (2024) reported experimental evidence of such over-reliance and highlighted the importance of appropriate trust calibration. Earlier work by Lee and See (2004) similarly distinguishes appropriate reliance from

simply increasing trust. This is particularly relevant to financial decision-making because users may act on recommendations that have meaningful personal consequences. The issue therefore extends beyond whether an AI system is trusted. It also concerns whether the interface helps users decide when to accept advice, when to examine it further, and when to disagree with it.

Trust remains an important part of the experience. Users need some reason to believe that an AI recommendation is sufficiently reliable to consider. Glikson and Woolley (2020), in their review of empirical research on human trust in AI, show that trust is shaped by multiple system and user factors rather than by a single interface feature. Recent research in AI-enabled financial services also indicates that trust can contribute to adoption. However, high trust is not automatically the same as good decision-making. A person may trust a system and still need to understand the reason behind its recommendation.

Autonomy and perceived control add another important dimension. A user may accept an AI recommendation while still feeling that they made the decision themselves. This distinction matters because recommendation acceptance alone cannot tell us whether the user felt involved in the process. Fink, Newman, and Haran (2024) found that increasing users' autonomy in recommendation systems can increase recommendation acceptance. Their work identified choice autonomy and control autonomy as mechanisms through which users can have a greater role in the recommendation process. For a UX study, this provides an important foundation for examining whether interface choices that preserve control can change how users respond to AI guidance.

The interface through which a recommendation is presented is therefore central to this research. Users do not experience an AI recommendation as an isolated sentence or prediction. They experience a complete screen containing headings, numbers, explanations, visual emphasis, choices, interaction controls, and actions. The order of information can affect what users notice first. The wording can make a suggestion sound optional or authoritative. The presence or absence of alternatives can affect whether users feel that they have a real choice. Research on explainable AI in finance also shows that transparency and interpretability are important concerns, but these qualities still have to be translated into an understandable interaction. An explanation can support understanding, but an explanation that is too long or unclear can create additional effort.

This research combines these ideas through a human-centered UX perspective. Rather than asking only whether users trust AI, the study asks how AI guidance should be presented so that it helps people make a decision without making the decision feel as though it has already been made for them. The study proposes prototype-based experimental conditions and simulated financial decisions to examine trust, perceived control, decision confidence, clarity, usability,and willingness to act. The focus is therefore not on designing a more persuasive AI recommendation, but on understanding how the interface can support informed human judgment. Literature Review

The literature relevant to this research is used to establish the variables and relationships tested in the study rather than to provide a standalone review of the field. Four areas are especially relevant: algorithmic recommendations and human decision-making, trust and appropriate reliance, autonomy and perceived control, and interface presentation and decision confidence. Together, these studies provide the theoretical basis for the experimental design. AI-generated recommendations and human decision-making

AI-based recommendation systems are increasingly used to support decisions rather than simply provide information. In a recommendation interaction, the user receives a system-generated suggestion and then has to decide what role that suggestion should play. The user may accept it, question it, compare it with another option, or reject it. This makes recommendation systems different from interfaces that only display facts. The recommendation becomes part of the user's reasoning process. Research on algorithmic judgment also shows that people can respond differently to algorithmic and human advice, which makes the user's relationship with the recommendation important to study.

Research by Fink, Newman, and Haran is particularly relevant because it examined how autonomy influences recommendation acceptance. Across three online experiments, the researchers found that users were more likely to accept algorithmic recommendations when they had greater choice autonomy, such as seeing multiple recommendations, and when they had greater control over the recommendation process. This suggests that acceptance does not depend only on the recommendation itself. The way users participate in the recommendation process can also shape their response.

This finding is useful for the present research because it shifts attention from the question of whether an AI recommendation is persuasive to the question of how the user is positioned around it. If an interface presents one recommendation as the

only obvious answer, the user may perceive less room for independent evaluation. If the interface presents the recommendation as one useful input among meaningful alternatives, the user may experience a different relationship with the system. The research therefore treats user participation as an important part of recommendation design.

The issue becomes more sensitive when recommendations concern financial decisions. Financial choices often involve personal goals, risk tolerance, time horizons, and circumstances that may not be completely represented in a system. A recommendation can therefore be useful while still being incomplete from the user's perspective. The interface should not make the user believe that the system has perfect knowledge of their situation. Instead, it should help the user understand what the recommendation is based on and what remains for the user to decide.

Recent work on recommendation explanations reinforces this point. Chen, Lin, and Yi (2025), studying conversational recommendations, found that post-hoc explanations can increase perceived transparency and interpretability and can subsequently influence trust. Their findings also show that explanation preferences can differ depending on decision context, interface type, and decision-making style. Although their work is not specifically about financial recommendations, it supports the broader idea that explanation should be considered as part of interface design rather than as an independent technical feature. Wang and Ding (2024) similarly show that the effects of explanation on trust and decision performance depend on how people process the explanation.

Taken together, the literature suggests that an AI recommendation is better understood as a social and interactive element within a decision process. The user is not simply receiving an answer. The user is interpreting the answer, comparing it with personal expectations, deciding how much weight to give it, and eventually taking or not taking an action. This makes recommendation presentation highly relevant to UX research.

Trust and reliance on AI recommendations

Trust is one of the most frequently discussed factors in human-AI decision-making. When users cannot fully inspect how an AI system reached a recommendation, they need to form some judgment about whether the system is reliable enough to use. Trust can develop from perceived competence, consistency, transparency, previous experience, and the way the system communicates its recommendation. In a financial context, the question becomes even more important because the recommendation may affect choices involving money, saving, spending, or other personally significant outcomes.

Klingbeil, Grützner, and Schreck (2024) provide important evidence for this research. Their experimental study found that people can over-rely on AI advice, including cases where AI advice conflicted with available contextual information or participants' own assessments. The study also found a relationship between trust in the advisor and reliance behavior. These findings show why increasing trust should not automatically be treated as the goal of AI interface design. A better goal is appropriate reliance, where users can make use of AI assistance while remaining able to recognize uncertainty and evaluate the recommendation themselves. This distinction is consistent with the earlier human-automation literature of Parasuraman and Riley (1997).

This distinction between trust and appropriate reliance is central to the present study. If an interface simply increases confidence in the AI without helping users understand its limitations, it may produce a strong but poorly calibrated response. A user might follow the recommendation because the system looks authoritative, because the interface highlights one answer, or because the wording makes disagreement feel unnecessary. Such a response would not necessarily demonstrate informed decision-making.

Research on AI financial service adoption also shows that trust can have multiple dimensions. Chang, Park, and Dinh (2026) examined different dimensions of trust in AI financial services and their relationship with adoption. This supports the relevance of trust in financial AI interactions, while also showing that trust is not a single simple feeling. Users may trust a system because they believe it is competent, because they believe it is acting appropriately, or for a combination of reasons.

For the present research, trust is therefore treated as one component of the decision experience. The study is interested not only in whether participants report trust, but in how trust relates to perceived control, confidence, clarity, and willingness to act. This approach allows the research to examine whether an interface can support confidence without encouraging blind reliance.

A practical UX question follows from this literature: how can a system help users know when to trust an AI recommendation and when to question it? The answer may involve the presentation of reasons, alternatives, uncertainty, user inputs, and opportunities for reflection. The interface can make these elements visible without overwhelming the user. This balance between support and independence is one of the central design problems explored in this study. User autonomy and perceived control

User autonomy refers to the ability to make a choice independently and to experience the choice as one's own. In an AI-assisted interface, autonomy does not necessarily mean that the user ignores the system. Instead, it can mean that the user has meaningful opportunities to decide how the recommendation should influence the final decision. Perceived control is closely related because it captures how much control the user feels they have during the interaction.

This distinction is important because the same final action can have different meanings. Two users may both accept an AI recommendation, but one may feel that the recommendation was useful evidence while the other may feel that the system effectively made the decision for them. A conventional usability measure such as task completion would treat these

interactions as similar. A user-centered study of AI decision-making needs to look deeper into how the decision was experienced.

Fink, Newman, and Haran's findings provide a useful foundation for this idea. Their work suggests that greater choice autonomy and control autonomy can increase recommendation acceptance. For the present research, the important implication is not simply that more choice is always better. Rather, the interface should provide meaningful control that helps users understand that the recommendation is support rather than an unavoidable instruction.

Perceived control can be influenced by small interface decisions. Showing more than one reasonable path can communicate choice. Allowing users to open an explanation can communicate that the recommendation can be examined. Making the user's own inputs visible can remind them that the recommendation is based partly on information they provided. Providing a clear way to change an assumption can make the system feel adjustable rather than fixed. These interactions may affect how much ownership the user feels over the final decision. The present study therefore treats control as an experience to be measured, not simply as the presence of a button or setting.

The research therefore connects interface design to perceived control. The proposed experimental conditions will examine different ways of presenting the recommendation and supporting information. The intention is not to assume in advance that one condition will always produce better outcomes. Instead, the study will compare how users describe their experience of control, trust, confidence, clarity, and willingness to act.

Interface presentation and decision confidence

Users do not interact with an AI recommendation in isolation. They encounter it through a screen or conversational interface that determines what is visible, what appears first, how information is grouped, and what actions are available. Amershi et al. (2019) place this kind of interaction design within the broader problem of making AI systems understandable and usable. The interface can therefore influence how much effort users spend understanding the recommendation and how confident they feel when moving toward a final decision.

Important presentation choices include the wording of the recommendation, the amount of supporting information, the visibility of alternatives, the location of explanations, the use of visual hierarchy, and the ease of questioning or changing the recommendation. A statement such as a direct instruction can create a different experience from a statement that describes the recommendation as a suggested option. Likewise, placing the explanation immediately next to the recommendation can create a different experience from hiding it several interactions away.

Recent work on recommendation explanations reinforces this point. Chen, Lin, and Yi (2025) found that post-hoc explanations can improve perceived transparency and interpretability in conversational recommendation settings. Their results also suggest that explanation preferences vary across decision contexts and interfaces. This supports a more flexible view of

explainability: the question is not simply whether an explanation exists, but whether the form of explanation fits what the user needs at that point in the decision.

More information is also not automatically better. A detailed explanation can increase cognitive effort if it is poorly structured or uses language that users do not understand. For a financial recommendation, an interface may need to communicate the key reason for the recommendation first, followed by additional detail for users who want to inspect it. This is consistent with a human-centered approach in which information is available without forcing every user to process the same level of detail.

Decision confidence is another important part of the experience. Confidence does not mean that the user is guaranteed to make the objectively correct choice. It refers to how certain the user feels that they understand the recommendation and can make a decision. Wang and Ding (2024) found that explanation can affect decision performance, while their results also show that self-reported trust does not necessarily move in the same way as behavioral outcomes. A useful interface may therefore support confidence by reducing unnecessary ambiguity while still making uncertainty visible when it matters.

These ideas lead to a central proposition of this research: the interface is not merely a container for the AI recommendation. It can become part of the decision itself. The design may influence what users understand, what they trust, how much control they feel, and how confidently they act.

Research gap

The existing literature provides evidence that AI recommendations, trust, autonomy, explanation, and reliance can influence human decision-making. Research is also increasingly examining AI-enabled financial services and the role of trust in their adoption. These studies establish a strong foundation for the present research, but they leave room for a focused examination of the interface-level decision experience. The gap identified here is therefore not the absence of research on trust or explainability; it is the opportunity to examine how these factors operate together during a specific user decision.

Many studies ask whether users trust AI, whether users accept an AI recommendation, whether explainability increases trust, whether autonomy changes acceptance, or whether people

over-rely on AI. These are important questions, but they often focus on one relationship at a time. From a UX perspective, the user experiences these factors together. The user sees a recommendation, interprets it, considers the explanation, evaluates alternatives, experiences a sense of control, forms an impression of trust, and then decides whether to act. The present research turns this sequence into an experimental interface problem.

The opportunity for this study is therefore to connect these elements within one user-centered framework. The proposed framework begins with the presentation of the AI recommendation and considers how presentation affects understanding and clarity, how those experiences relate to trust and perceived control, and how trust and control may contribute to decision confidence and willingness to act. The framework does not assume a simple one-way relationship; rather, it recognizes that these factors can influence each other during the interaction.

This research gap is particularly relevant to financial UX because the cost of a poorly understood recommendation can be meaningful to the user. The study therefore focuses on the design of the decision experience rather than attempting to evaluate the financial quality of a particular investment or product recommendation.

Research Objectives

The main purpose of this research is to understand how the presentation of AI-generated financial recommendations influences the way users experience and respond to the decision. The study focuses on the interaction between recommendation presentation, trust, perceived control, decision confidence, and willingness to act. These variables were selected because prior human-AI research connects recommendation acceptance with autonomy and connects reliance with trust calibration.

The research seeks to examine whether different interface presentations change how clearly users understand an AI recommendation. Clarity is important because a recommendation cannot meaningfully support a decision if the user cannot identify what the system is suggesting, why it is suggesting it, or what information remains uncertain. Human-AI interaction guidelines from Amershi et al. (2019) provide a useful design basis for treating clarity and user understanding as interaction concerns rather than purely technical properties.

The study also seeks to understand the role of perceived control. It examines whether interfaces that provide meaningful choice, visible reasoning, and opportunities to question or explore the recommendation create a stronger feeling that the user remains responsible for the final decision. This objective follows the autonomy findings reported by Fink et al. (2024), but the present study tests control at the interface level.

Another objective is to examine the relationship between trust and decision behavior. The study does not treat high trust as automatically desirable. Instead, it explores whether trust is associated with willingness to act and whether this relationship changes when users experience different levels of control and clarity.

Finally, the research aims to translate the findings into practical UX implications. The intended outcome is a set of design considerations for AI-assisted financial interfaces that help users understand recommendations, maintain involvement, and make decisions with an appropriate level of confidence.

Research Questions

The central research question is: How does the presentation of AI-generated financial recommendations influence users' perceived control, decision confidence, trust, and willingness to act on the recommendation?

A supporting question asks how the presentation of an AI-generated financial recommendation influences users' understanding and trust in that recommendation. This considers both the information itself and the interface choices that communicate it.

A second supporting question asks how interface presentation affects users' perceived control over the final financial decision. This focuses on whether users feel they are making the decision with AI assistance or simply following a decision that appears to have already been made for them.

A third supporting question asks how trust and perceived control relate to users' decision confidence and willingness to act on an AI-generated recommendation. This allows the study to examine the relationships among the psychological and behavioral measures rather than treating each one as an isolated outcome.

Conceptual Framework

The proposed conceptual framework connects AI recommendation presentation with the user's experience of making a decision. The first part of the framework is recommendation presentation. This includes how the recommendation is worded, how much supporting information is shown, whether alternatives are visible, whether an explanation can be accessed, and how much opportunity the user has to question or modify the recommendation. These dimensions are informed by research on explainability, recommendation interfaces, and human-AI interaction.

The next part of the framework is understanding and clarity. Before users can evaluate an AI recommendation, they need to understand what the system is suggesting and what information supports the suggestion. Clarity is therefore treated as an important experience that can influence later

judgments. If users cannot understand the recommendation, trust and confidence may be based more on presentation than on actual understanding.

Trust and perceived control form the central user-experience layer of the framework. Trust concerns whether the user considers the AI recommendation sufficiently reliable to use. Perceived control concerns whether the user feels that the final decision remains under their control. These two experiences may interact. A user can trust an AI system while still wanting control, and a user can feel control while remaining uncertain about the recommendation.

Decision confidence is positioned as an outcome of the decision experience. When users understand the recommendation and feel able to evaluate it, they may feel more prepared to make a final choice. However, confidence should not be interpreted as proof that the decision is correct. In this study, it represents the user's confidence in their own decision process.

Willingness to act represents the behavioral intention to follow or use the recommendation. The framework therefore moves from how the recommendation is presented, through how the user understands and evaluates it, toward how confidently the user decides whether to act. The framework is deliberately human-centered because the final stage remains a user decision rather than an automatic system action.

Research Hypotheses

The first proposed hypothesis is that interface presentation that provides clearer recommendation information and accessible supporting reasoning will be associated with higher perceived clarity than a presentation that provides less contextual support. This hypothesis is grounded in research showing that explanations and interface presentation can influence understanding and transparency.

The second proposed hypothesis is that interfaces providing meaningful user choice and visible opportunities to question or inspect the recommendation will be associated with higher perceived control. This hypothesis follows from the autonomy findings reported by Fink et al. (2024), while extending the question to a specific financial decision interface.

The third proposed hypothesis is that higher perceived clarity will be associated with greater trust in the AI recommendation, while recognizing that trust may also depend on other factors such as perceived competence and the user's prior expectations.

The fourth proposed hypothesis is that higher perceived control will be associated with greater decision confidence because users may feel more ownership of the final choice when they can actively evaluate the recommendation. This is treated as a testable relationship rather than an assumed outcome.

The fifth proposed hypothesis is that trust and perceived control will both be related to willingness to act on the AI recommendation, but that perceived control may help distinguish informed acceptance from passive reliance.

These hypotheses are proposed relationships for empirical testing. They are not presented as findings. The actual direction and strength of relationships will depend on the data collected from participants.

The study will use a user-centered experimental design in which participants interact with prototype interfaces that present the same underlying financial recommendation in different ways. The purpose is to isolate the effect of presentation rather than compare different financial products. This approach follows the broader human-AI research practice of studying how interface conditions affect user behavior and perception while keeping the underlying task comparable.

The study uses a user-centered experimental approach. The main method is to create prototype digital interfaces that present AI-generated financial recommendations in different ways and ask participants to complete simulated financial decision-making tasks. The purpose of the experiment is to compare experiences rather than to evaluate a real financial product or provide real financial advice.

A prototype-based method is appropriate for this research because the central variables are closely connected to interface presentation. A purely descriptive survey may identify whether users trust AI, but it would be less suitable for testing how changes in wording, explanation, alternatives, and control affect that experience. By presenting participants with controlled interface conditions, the research can compare responses while keeping the underlying scenario as consistent as possible.

The study is designed around the idea that the recommendation itself should remain conceptually similar across conditions while the presentation changes. This makes it possible to focus on the user experience surrounding the recommendation. The experimental conditions can vary in factors such as the visibility of explanation, availability of alternatives, wording of the recommendation, and opportunities for the user to inspect or adjust information.

The research is exploratory and comparative. It is not intended to establish that one interface will always be better for every financial decision. Financial decision-making varies by context, experience, risk, and personal goals. Instead, the study aims to identify patterns that can help designers understand how interface presentation may support or reduce user agency.

Experimental Interface Conditions

The proposed experiment can compare interface conditions that differ in the amount and form of decision support provided. One condition can present the AI recommendation in a relatively direct form, where the recommended action is visually prominent and limited supporting context is immediately visible. This condition represents an interface in

which the recommendation is easy to identify but may provide fewer opportunities for reflection.

A second condition can provide the same recommendation with a concise explanation of the main factors that influenced it. The explanation should be understandable without requiring specialist knowledge. The purpose is to examine whether accessible reasoning changes clarity, trust, perceived control, or confidence.

A third condition can provide the recommendation together with meaningful alternatives or an explicit opportunity to explore other options. This condition is informed by research on choice autonomy. The intention is to examine whether giving users more control over the recommendation space changes their experience of ownership and willingness to act.

A further condition, if resources allow, can combine explanation with user control. In this condition, the user can see the recommendation, understand the key reasons, and explore alternatives before deciding. The purpose is not to assume that combining features will automatically produce the strongest response, but to examine whether the combination creates a different decision experience from a more directive presentation.

All conditions should maintain consistent visual quality, terminology, scenario information, and interaction length as far as practical. Only the intended interface differences should change. This helps reduce the possibility that participants respond simply because one prototype looks more polished or contains more decorative elements.

Experimental Scenario

The study will use a simulated financial decision scenario rather than asking participants to make a real financial transaction. The scenario should be understandable to participants without requiring professional financial knowledge. Participants will be given a realistic personal context, a defined goal, and a limited set of information. The AI system will then provide a recommendation based on the scenario.

The scenario should create a genuine need to evaluate the recommendation rather than making the correct response obvious. For example, participants may be asked to consider how they would allocate a limited amount of available money toward competing goals. The exact financial values can be adjusted during pilot testing so that the task is neither too easy nor unnecessarily complex. The purpose is to create a decision in which participants can reasonably agree or disagree with the AI recommendation.

Participants will be told that the scenario is simulated and that the AI recommendation is part of the research prototype. They will be asked to behave as they would when evaluating a recommendation in a digital financial interface. After reviewing the recommendation and available information, they will make a final decision and answer questions about their experience.

The scenario is designed to study decision behavior and perception, not to test participants' financial knowledge. Therefore, the questionnaire should avoid questions that reward specialist terminology. The research is interested in whether participants understand the recommendation, feel in control, trust the system appropriately, and feel confident about their decision.

Participants

The study is intended for adult participants who regularly use digital interfaces and have experience making everyday financial decisions. The participant group does not need to consist of professional financial experts because the research focuses on ordinary user interaction with AI-generated guidance. Participants should be able to understand the study instructions and interact with the prototype independently.

Participant recruitment should aim for variation in age, educational background, digital experience, and familiarity with AI tools where practical. Such variation can help the study avoid treating one narrow user group as representative of all users. At the same time, the study should record relevant background characteristics so that differences in familiarity with AI or financial interfaces can be considered during analysis.

Because the study is proposed rather than completed, the final sample size should be determined according to the available research setting, time, and statistical requirements. A pilot study can first be used to identify usability problems and estimate completion time. The final report should state the actual number of participants, recruitment method, demographic characteristics, and exclusions after data collection rather than presenting these details as completed results in advance.

Study Procedure

The study begins with an introduction explaining that participants will interact with a prototype AI-assisted financial interface. Participants will be informed that the system is simulated for research purposes and that they should evaluate the recommendation rather than assume it is automatically correct. This instruction is important because the study is examining decision-making rather than obedience to the system.

Participants will then receive the financial scenario and the information needed to understand the task. They will be given enough time to read the scenario and interact with the prototype. Depending on the experimental design, each participant may experience one interface condition or several

counterbalanced conditions. If several conditions are used, the order should be varied where possible to reduce learning and order effects.

After reviewing the recommendation, participants will make a final decision through the prototype. The system will record the selected action and relevant interaction data if the prototype supports this. Immediately after the task, participants will answer questions about clarity, trust, perceived control, confidence, usability, and willingness to act.

A short qualitative response can also be collected by asking participants to explain why they accepted, questioned, or rejected the recommendation. This open response can help identify reasons that a numerical scale may not capture. For example, a participant may report that they accepted the recommendation because the explanation matched their personal goal, or they may report that they rejected it because the interface felt too certain.

The study concludes with a debriefing. Participants should be reminded that the recommendation was simulated and that no real financial decision was being made. Any personal information collected for research should be minimized and handled according to the applicable institutional requirements.

Perceived control will be measured through statements about whether participants felt they could evaluate the recommendation, make their own choice, and change or question the information presented to them. The measure is conceptually related to the autonomy construct examined by Fink et al. (2024).

Trust will be measured through statements that ask participants how reliable, credible, or dependable they considered the AI recommendation. The wording should avoid suggesting that the participant ought to trust the system. This is important because the research aims to measure trust as an experience rather than encourage it during the study.

Decision confidence will be measured through statements about whether participants felt prepared to make the decision, understood the important information, and were confident in their final choice. Confidence will be treated as a subjective judgment about the decision process rather than evidence that the participant made a objectively correct financial choice.

Decision confidence will be measured through participants' reported confidence in their final decision and in their understanding of the recommendation. The measure should distinguish confidence in the decision process from confidence that the financial outcome will definitely be successful.

Clarity will capture whether participants understood what the AI recommended and why. This can include questions about how easy the recommendation was to interpret and whether the supporting information was sufficient. Usability will capture the ease of navigating and interacting with the prototype.

Willingness to act will capture whether participants would consider following the recommendation in a similar situation. This measure should be interpreted carefully because willingness to act is not the same as proof that the recommendation is financially appropriate.

A short open-ended question can capture what participants considered most important before making their final decision. This qualitative information can be used to interpret patterns in the quantitative responses.

Questionnaire Design

The questionnaire should use simple statements and a consistent response scale so that participants can focus on the experience rather than understanding complicated measurement language. A five-point agreement scale can be used, ranging from strong disagreement to strong agreement. The exact scale should remain consistent across the main constructs unless there is a methodological reason to change it.

Trust items can ask whether the participant considered the AI recommendation reliable, whether they believed the recommendation was based on useful information, and whether they felt comfortable considering it. Perceived control items can ask whether the participant felt that the final decision remained their own, whether they could question the recommendation, and whether they felt free to choose a different option.

Decision confidence items can ask whether participants felt confident about their final choice and whether they felt they had enough information to decide. Clarity items can ask whether the recommendation was easy to understand and whether the main reason for the recommendation was clear. Usability items can ask whether the interface was easy to navigate and whether the interaction felt straightforward.

Willingness-to-act items can ask whether the participant would consider following the recommendation, whether they would use a similar AI recommendation in the future, and whether they would recommend using such an interface for similar decisions. These items should be analyzed alongside perceived control rather than alone.

Because the research is specifically concerned with user agency, the analysis should avoid treating the highest acceptance rate as automatic evidence of better design. A condition could produce more recommendation acceptance because it is more persuasive, not because users understand it better. The analysis will therefore examine acceptance together with perceived control, clarity, trust, and decision confidence.

The quantitative analysis will begin with descriptive statistics to summarize participants' responses across the main

constructs. Mean scores, response distributions, and variation can be used to understand the overall pattern of trust, perceived control, confidence, clarity, usability, and willingness to act. These summaries should be reported separately for each experimental interface condition.

If the final sample size and design support inferential testing, comparisons between interface conditions can be conducted using an appropriate statistical method. The specific test should depend on whether participants experience one condition or multiple conditions and on the distribution of the collected data. The final paper should report the actual test used rather than selecting a method only after seeing which produces a preferred result.

Relationships among trust, perceived control, decision confidence, and willingness to act can also be examined using correlation or regression-based analysis where appropriate. This can help identify whether users who report greater control also report greater confidence or willingness to act. Such relationships should be interpreted as associations unless the experimental design and analysis justify stronger causal conclusions.

Open-ended responses can be analyzed thematically. Responses can be grouped around themes such as understanding the recommendation, desire for more information, perceived pressure, usefulness of alternatives, trust in the system, uncertainty, or feeling responsible for the final choice. These themes can help explain why numerical responses differ across conditions.

The study may reveal that a highly persuasive presentation produces strong willingness to act while also reducing perceived control. If that pattern appears, it would be important because it would distinguish recommendation effectiveness from user-centered decision support. The study may also find that a condition with more explanation produces better understanding but lower ease of use. Such results would suggest a trade-off rather than a single universally optimal interface.

Expected Findings and Interpretation

The present paper does not report completed empirical findings because the proposed study has not yet been conducted. The following section therefore describes the patterns that the research is designed to examine rather than claiming results. This distinction is important because expected findings must not be presented as observed evidence.

One expected possibility is that interfaces providing clear explanations and meaningful user control will produce stronger perceptions of involvement than interfaces that present a recommendation in a highly directive way. Participants may report that they feel more responsible for the final decision when they can see why the recommendation was made and can explore alternatives.

Another possibility is that trust and perceived control will not always move together. A user may trust the AI but still feel that the interface is too directive. Conversely, a user may feel highly in control but remain uncertain about whether the recommendation is reliable. Such patterns would be useful because they would show that trust and agency are related but distinct aspects of the decision experience.

The study may also reveal that additional explanation has a point of diminishing value. Participants could find a concise explanation useful while finding a long explanation difficult to process. If this occurs, the design implication would not be to remove explanations, but to structure them progressively so that the most important information is available first and deeper detail can be accessed when needed.

Another possible finding is that users may accept recommendations more readily when alternatives are visible. This would be consistent with research on choice autonomy, but the present study would examine the effect specifically through an interface and financial decision context. The important point would be whether alternatives help users feel more capable of evaluating the recommendation rather than simply giving them more options to click.

The final analysis should therefore remain open to different patterns. The value of the study lies in identifying how presentation, control, trust, and confidence interact rather than in proving that one interface element is universally superior.

UX Design Implications

The first design implication is that an AI recommendation should be presented as decision support rather than as an unquestionable instruction. The visual hierarchy can still make the recommendation easy to find, but the interface should communicate that the user remains the decision-maker. Small wording choices can contribute to this distinction by describing the output as a recommendation, suggestion, or explanation of an option rather than presenting it as an absolute command.

The second implication concerns explanation. Users should be able to understand the main reason behind a recommendation without needing to read a long technical explanation. A concise summary can communicate the key factors, while a secondary interaction can reveal more detail for users who want to inspect the reasoning. This approach can support different levels of user interest without placing the same cognitive burden on everyone.

The third implication is meaningful choice. If the system can identify several reasonable options, showing only one may unintentionally make the interface feel more authoritative than necessary. Providing alternatives can help communicate that the AI is assisting with the decision rather than owning it. However, alternatives should be meaningful and

understandable; simply adding more options can increase confusion if users cannot distinguish between them.

The fourth implication is visible user input. When a recommendation depends on information about the user's goals or preferences, the interface can make the important inputs visible. This helps users recognize the connection between what they told the system and what the system recommends. It can also give them a clear opportunity to correct information that does not represent their situation.

The fifth implication is appropriate uncertainty. An interface should avoid creating a stronger impression of certainty than the underlying recommendation supports. This does not mean making every screen full of warnings. Instead, uncertainty can be communicated in simple language where it matters, especially when the recommendation depends on assumptions or incomplete information.

The final implication is that acceptance should not be the only UX success metric. A design that increases recommendation acceptance may appear effective, but if users feel unable to question the system, it may also create over-reliance. Designers should therefore consider trust, perceived control, understanding, confidence, and willingness to act together.

Ethical Considerations

Because the study concerns financial decision-making and AI recommendations, ethical communication is important even though the tasks are simulated. Participants should understand that the prototype does not provide real financial advice and that they should not use the research interface to make actual financial decisions. This prevents confusion between the research scenario and professional financial guidance.

Participant autonomy should also be respected throughout the study. Participants should be free to withdraw according to the research setting's informed-consent requirements. The study should avoid collecting unnecessary personal financial information. If demographic or experience data are collected, they should be limited to information necessary for the research question.

The design itself should not intentionally pressure participants into accepting a recommendation. Since the research is about perceived control, a coercive or manipulative interface would introduce an ethical problem as well as a methodological one. The experimental conditions should differ in presentation while maintaining respect for the participant's ability to choose.

Finally, the interpretation of the research should avoid presenting AI as either inherently trustworthy or inherently harmful. The purpose is to understand how people interact with AI recommendations and how interface design can support better human involvement. The findings should be discussed with appropriate uncertainty and should not be generalized beyond the participant group and study conditions.

Validity, Reliability, and Pilot Testing

Before the main study, a pilot test should be conducted with a small group of participants to identify unclear wording, confusing interface elements, technical problems, and unrealistic task demands. Pilot participants can be asked to explain what they think the AI recommendation means and what they believe they are expected to do. This can reveal whether the prototype communicates the intended distinction between recommendation and decision.

The experimental conditions should be checked for visual equivalence. If one condition contains substantially more polished visuals, larger typography, or stronger color emphasis, participants may respond to the visual quality rather than the intended recommendation presentation. Keeping layout quality and interaction patterns consistent helps improve internal validity.

Questionnaire items should be reviewed for clarity and redundancy. If several questions appear to ask exactly the same thing, participants may experience unnecessary repetition. At the same time, each construct should have enough items to represent it meaningfully. Reliability analysis can be conducted after data collection if the sample size and measurement design support it.

The final study should also report limitations related to the prototype. Participants interacting with a simulated interface may behave differently from people making real financial decisions with real consequences. The study can still provide valuable UX evidence, but the findings should be interpreted within the boundaries of the experimental setting.

For UX designers, the practical challenge is therefore one of balance. Too little information can make the recommendation difficult to evaluate, while too much information can create unnecessary cognitive effort. Research on explainability and recommendation interfaces suggests that explanation should be understandable and relevant to the decision context rather than added simply for transparency's sake. The findings of this study can translate that principle into concrete interface choices.

The central idea of this research is that AI-assisted decision-making should not be understood only as a question of whether users follow recommendations. Following a recommendation is the final visible action, but many experiences occur before that action. The user first notices the recommendation, interprets it, decides whether it seems credible, considers personal goals, evaluates the available alternatives, and experiences some degree of control or pressure. Interface design participates in each of these stages.

This perspective also helps explain why trust alone is not sufficient as a design goal. A system can be trusted because it appears competent, familiar, or authoritative. Yet a user may still need opportunities to question it. If trust becomes a shortcut that prevents evaluation, it can contribute to over-reliance. If trust is supported by understandable information and meaningful control, it may instead become part of a more balanced decision process.
Perceived control is therefore important because it adds a human dimension that acceptance measures can miss. When users feel that the decision remains theirs, they may be more likely to treat the AI as a support tool rather than a replacement for their own judgment. This does not mean that users should ignore AI recommendations. It means that the interface should make room for active evaluation.

The research also highlights the psychology of interface elements. A recommendation is not interpreted only through its words. Position, hierarchy, emphasis, interaction cost, and availability of alternatives can all influence how users understand the relationship between themselves and the system. A visually dominant recommendation with little opportunity to explore may communicate a different level of authority from a recommendation presented alongside reasons and alternatives.

The interface can also communicate authority through visual hierarchy and language. A large, visually dominant recommendation with a single action may feel more final than a recommendation presented with reasons, alternatives, and a clear opportunity to reconsider. This does not mean that every interface should add more options. Instead, the research asks whether the available interaction makes the user's role in the final decision visible.

The proposed study contributes to this discussion by bringing these concerns into one experimental interface context. The study does not attempt to replace research on AI accuracy, financial outcomes, or algorithmic fairness. Instead, it focuses on the layer where users meet the

AI: the interface through which a recommendation becomes understandable, questionable, and actionable.

Limitations

trustworthy,

The first limitation is that the study uses simulated financial decision-making rather than real financial transactions. Participants may therefore experience less emotional or practical pressure than they would in a real situation. The results should be interpreted as evidence about interface experience and decision perception rather than proof of real-world financial behavior.

A second limitation is the participant sample. If the study is conducted with a relatively small or convenience-based sample, the findings may not represent all users. Differences in age, financial experience, digital literacy, and previous exposure to AI may influence responses. These factors should therefore be reported clearly and considered when interpreting the results.

A third limitation concerns the prototype itself. A research prototype cannot reproduce every feature of a fully developed AI financial system. Participants may respond differently when recommendations are connected to real personal data, real-time information, or real financial consequences. Future research could extend the study using more realistic settings while maintaining appropriate ethical protections.

A further limitation is that self-reported measures such as trust and confidence describe participants' perceptions and may not perfectly predict behavior. Combining questionnaire responses with observed interaction behavior and qualitative explanations can provide a richer understanding, but even these measures cannot capture every aspect of a complex decision.

Finally, the research focuses primarily on interface presentation and user agency. Other important issues such as model accuracy, data quality, fairness, privacy, regulation, and financial literacy remain relevant but are not the main focus of this study. They should be considered when applying the findings to real systems.

Future Scope

Future research can extend this work by testing AI recommendation interfaces with larger and more diverse participant groups. Different levels of financial experience and AI familiarity may influence how users interpret recommendations and how much control they expect from an interface. Comparing these groups could help identify where particular design patterns are most useful.

Future studies could also examine more realistic decision contexts, including longer-term interactions in which users return to the same system repeatedly. Trust and perceived control may change with experience. A user who initially questions every recommendation may develop a different interaction pattern after repeated successful experiences, while a user may also become overly dependent on a system that has previously appeared reliable.

Another area for future research is adaptive explanation. Instead of presenting the same explanation to every user, an interface could allow users to choose the level of detail they need. Future experiments could compare concise explanations, progressive disclosure, visual evidence, source information, and interactive questioning to understand which approaches support appropriate trust and confidence.

Future research could also examine the relationship between recommendation framing and action. The same AI output may produce different responses depending on whether it is framed as a suggestion, a prediction, a warning, or a recommended option. Testing these differences can help designers understand how language and visual hierarchy shape the psychological experience of AI assistance.

Finally, future work can connect the UX findings with technical and ethical research. Interface design alone cannot solve problems caused by inaccurate data, biased models, poor recommendations, or inappropriate system objectives. A more complete human-centered AI approach would consider the quality of the underlying recommendation together with how that recommendation is communicated to the person who must decide.

This research examines a simple but important question: when AI recommends what people should do with their money, who still feels that they are making the decision? The question becomes increasingly relevant as AI systems move from providing information toward generating personalized guidance. The value of such systems depends not only on their technical capability but also on how people understand and use their recommendations. Prior work on trust, reliance, autonomy, and explainability provides the foundation for this study, but the present research places those ideas inside a specific interface and decision-making experiment.

The literature reviewed in this paper shows that trust, autonomy, recommendation acceptance, explanation, and reliance are connected to human-AI decision-making. Research has also demonstrated the possibility of over-reliance and has shown that giving users greater autonomy can affect recommendation acceptance. In financial AI contexts, trust is an important factor in adoption. At the same time, explanation research suggests that how information is presented can influence transparency, interpretability, and decision performance. The present study uses these findings as a basis for testing interface-level relationships rather than treating them as conclusions about the proposed prototype.

Building on this evidence, the present research takes an interface-centered perspective. It proposes that AI recommendation presentation can influence clarity, perceived control, trust, decision confidence, and willingness to act. The study therefore moves beyond the question of whether users trust AI and focuses on what happens between receiving an AI recommendation and making the final choice. The experimental method is designed to test these relationships rather than assume them.

The proposed experimental method is designed to compare different prototype presentations of AI-generated financial recommendations. By examining the user experience across conditions, the research aims to understand whether explanation, alternatives, and meaningful control can help users engage with AI without feeling that the system has taken over the decision. The findings, once the study is completed, can provide evidence for practical UX decisions around recommendation wording, hierarchy, explanation, alternatives, uncertainty, and user control.

The central design principle emerging from this research is that useful AI assistance should support human judgment rather than quietly replace it. A good AI-assisted interface does not need to make users accept a recommendation immediately. It should help them understand what the system is suggesting, why it is suggesting it, what alternatives exist, and where their own judgment still matters. In that sense, the success of an AI recommendation interface should be measured not only by whether users follow the recommendation, but also by whether they understand the decision and feel that it remains meaningfully theirs.

The central design principle emerging from this research is that useful AI assistance should support human judgment rather than quietly replace it. A good AI-assisted interface does not need to make users accept a recommendation immediately. It should help them understand what the system is suggesting, why it is suggesting it, what alternatives exist, and where their own judgment still matters. In that sense, the success of an AI recommendation interface should be measured not only by whether users follow the recommendation, but also by whether they understand the decision and feel that it remains meaningfully theirs.

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 1–13. https://doi.org/10.1145/3290605.3300233

Buçinca, Z., Malaya, M. B., & Gajos, K. Z. (2021). To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, 1–14. https://doi.org/10.1145/3411764.3445773Cardillo, G., & Chiappini, H. (2024). Robo-advisors: A systematic literature review. Finance Research Letters, 62, 105119. https://doi.org/10.1016/j.frl.2024.105119

Chang, C.-Y., Park, S., & Dinh, M. C. (2026). Advice by algorithms: Dual dimensions of trust in AI financial service adoption. Journal of Retailing and Consumer Services, 88, 104537. https://doi.org/10.1016/j.jretconser.2025.104537

Chen, Q. Q., Lin, L. M., & Yi, Y. (2025). Tailoring explanations in conversational recommendations: The impact of decision contexts and user interfaces. Journal of Retailing and Consumer Services, 85, https://doi.org/10.1016/j.jretconser.2025.104281 104281.

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

Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv. https://arxiv.org/abs/1702.08608

Fahnenstich, H., Rieger, T., & Roesler, E. (2024). Trusting under risk—Comparing human to AI decision support agents. Computers in Human Behavior, 153, 108107. https://doi.org/10.1016/j.chb.2023.108107

Fatima, S., & Chakraborty, M. (2024). Adoption of artificial intelligence in financial services: The case of robo-advisors in India. IIMB Management Review, 36(2), 113–125. https://doi.org/10.1016/j.iimb.2024.04.002

Ferrario, A. (2024). Justifying our credences in the trustworthiness of AI systems: A reliabilistic approach. Science and Engineering Ethics, 30, 55. https://doi.org/10.1007/s11948-024-00522-z

Fink, L., Newman, L., & Haran, U. (2024). Let me decide: Increasing user autonomy increases recommendation acceptance. Computers in Human Behavior, 156, 108244. https://doi.org/10.1016/j.chb.2024.108244

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. Academy of Management Annals, 14(2), 627–660. https://doi.org/10.5465/annals.2018.0057

Khan, F. S., Mazhar, S. S., Mazhar, K., AlSaleh, D. A., & Mazhar, A. (2025). Model-agnostic explainable artificial intelligence methods in finance: A systematic review, recent developments, limitations, challenges and future directions. Artificial Intelligence Review, 58, 232. https://doi.org/10.1007/s10462-025-11215-9

Klingbeil, A., Grützner, C., & Schreck, P. (2024). Trust and reliance on AI—An experimental study on the extent and costs of overreliance on AI. Computers in Human Behavior, 160, 108352. https://doi.org/10.1016/j.chb.2024.108352

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. Human Factors, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392

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

Luo, H., Liu, X., Lv, X., Hu, Y., & Ahmad, A. J. (2024). Investors' willingness to use robo-advisors: Extrapolating influencing factors based on the fiduciary duty of investment advisors. International Review of Economics & Finance, 94, 103411. https://doi.org/10.1016/j.iref.2024.103411

Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. Academy of Management https://doi.org/10.5465/AMR.1995.9508080335Review, 20(3),709–734.

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. Human Factors, 39(2), 230–253. https://doi.org/10.1518/001872097778543886

Poursabzi-Sangdeh, F., Goldstein, D. G., Hofman, J. M., Vaughan, J. W., & Wallach, H. (2021). Manipulating and measuring model interpretability. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, https://doi.org/10.1145/3411764.34453151–52.

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1135–1144. https://doi.org/10.1145/2939672.2939778

Schoorman, F. D., Mayer, R. C., & Davis, J. H. (2007). An integrative model of organizational trust: Past, present, and future. Academy of Management Review, 32(2), 344–354. https://doi.org/10.5465/AMR.2007.24348410

Sheikh, A. A., Goel, S., & Kumar, Y. (2024). AI as decision aid or delegated agent: The effects of trust dimensions on the adoption of AI digital agents. Computers in Human Behavior: Artificial Humans, 2(2), 100094. https://doi.org/10.1016/j.chbah.2024.100094

Sullivan, V., & Weger, K. (2025). Transparency and explainability in AI-assisted decision making: Effects on trust, perceived reliability, confidence, and ease of understanding. Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 69(1). https://doi.org/10.1177/10711813251369473

Wang, P., & Ding, H. (2024). The rationality of explanation or human capacity? Understanding the impact of explainable artificial intelligence on human-AI trust and decision performance. Information Processing & Management, 61(4), 103732. https://doi.org/10.1016/j.ipm.2024.103732

Weber, P., Carl, K. V., & Hinz, O. (2024). Applications of explainable artificial intelligence in finance—A systematic review of finance, information systems, and computer science literature. Management Review Quarterly, 74, 867–907. https://doi.org/10.1007/s11301-023-00320-0

Yeo, W. J., Van Der Heever, W., Mao, R., Cambria, E., Satapathy, R., & Mengaldo, G. (2025). A comprehensive review on financial explainable AI. Artificial Intelligence Review, 58, 189. https://doi.org/10.1007/s10462-024-11077-7

Zhu, H., Vigren, O., & Söderberg, I.-L. (2024). Implementing artificial intelligence empowered financial advisory services: A literature review and critical research agenda. Journal of Business Research, 174, 114494. https://doi.org/10.1016/j.jbusres.2023.114494

Černevičienė, J., & Kabašinskas, A. (2024). Explainable artificial intelligence (XAI) in finance: A systematic literature review. Artificial Intelligence Review, 57, 216. https://doi.org/10.1007/s10462-024-10854-8

References

Appendix: Proposed Participant Questionnaire

The proposed questionnaire begins with a short instruction asking participants to answer based on the interface they have just used. The questionnaire is intended to measure experience rather than financial knowledge.

Trust can be assessed through statements such as: I considered the AI recommendation reliable; I believed the recommendation was based on useful information; I felt comfortable considering the recommendation; and I believed the system provided a reasonable recommendation for the scenario.

Perceived control can be assessed through statements such as: I felt that the final decision was mine; I felt free to disagree with the AI recommendation; I felt able to evaluate the recommendation before deciding; and I felt that I had enough control over the decision process.

Decision confidence can be assessed through statements such as: I felt confident about my final decision; I understood the recommendation well enough to decide; I felt prepared to make the decision; and I was confident that I had considered the important information.

Clarity and usability can be assessed through statements such as: The recommendation was easy to understand; the reason for the recommendation was clear; the information was organized well; the interface was easy to navigate; and the interaction felt straightforward.

Willingness to act can be assessed through statements such as: I would consider following this recommendation in a similar situation; I would consider using this type of AI assistance again; and I would feel comfortable using the interface as a decision-support tool.

The questionnaire can conclude with an open-ended question asking: What was the main thing that influenced your final decision? A second optional question can ask: Was there anything about the AI recommendation or interface that made you want to question the recommendation? These responses can provide qualitative evidence about user agency and trust.