Belief Bias, Source Preference, and Preference for Robo-Advisor*

Paul H. Y. Cheung Haomin He King King Li

November 7, 2026

Abstract

Do people prefer robo-advisors because they believe the algorithm is more accurate, or because they favor it independent of performance? This paper investigates preferences for human versus robo-advisors using a dynamic source preference framework that allows for belief bias. We use a two-part experiment in which advice is generated by real-life investors and ChatGPT, and subjects value advice before and after learning advisors' objective accuracy. This design enables a dynamic decomposition of belief bias from source preference. Before disclosure, subjects are more likely to prefer robo-advice over human advice. After learning that the two advisors perform similarly, many subjects revise their valuations away from the robo-advisor and toward the human advisor, indicating that initial robo-preference partly reflects biased beliefs about algorithmic competence. Yet, a substantial subset continues to prefer the robo-advisor after disclosure, revealing source preference independent of performance. We also find that subjects have a stronger preference for robo-advisor under uncertainty than under risk. The preference for robo-advisor is reduced significantly after the uncertainty on performance is resolved via the disclosure of accuracy information, supporting the competence hypothesis of Heath and Tversky (1991). Overall, the results show that belief bias and source preference jointly shape demand for robo-advice, and the preference for robo-advisor was stronger under uncertainty.

Keywords: Robo-Advisor, Belief Bias, Source Preference, Uncertainty, Experiment

JEL: C91, G4

*Cheung: Naveen Jindal School of Management, University of Texas at Dallas. Email: paul.cheung@utdallas.edu. He: Rady School of Management, UC San Diego. Email: hah076@ucsd.edu. Li: School of Business, Department of Economics and Finance, The Hang Seng University of Hong Kong. Email: likingking@gmail.com. We thank Soo Hong Chew, Songfa Zhong, Lingbo Huang, Shuguang Jiang, Ernest Lai, Jaimie Lien, Arno Riedl, Liyan Yang, Jie Zheng, conference participants at D-TEA 2026, FUR Conference 2026, China Behavioral and Experimental Economics Forum 2025, and SZU Public Policy and Experimental Economics Workshop 2023 for their helpful comments and discussions. The experimental protocol was approved by the IRB at UTD (IRB-25-276), and pre-registered on AsPredicted (#238,320). Li gratefully acknowledges financial supports from the National Natural Science Foundation of China under grant No. 71973099.

1 Introduction

Robo-advisors have emerged as one of the most disruptive developments in asset and wealth management (Shanmuganathan, 2020), attracting substantial attention from financial institutions, investors, and regulators. A growing number of firms have entered the robo-advisory market, with prominent platforms such as Vanguard Digital Advisor, Schwab Intelligent Portfolios, and Betterment managing substantial assets. For instance, Vanguard Digital Advisor currently oversees more than $340 billion in assets and serves about 800,000 clients, as reported in the October 2026 Form ADV filed with the United States Securities and Exchange Commission (SEC).[1] The increasing use of robo-advisors raises an important question: what drives individuals' preferences for robo-advice relative to advice from human advisors?

More broadly, algorithmic advice is spreading well beyond investing, from AI tools that assist medical diagnosis (Schwalbe and Wahl, 2020) and research for legal cases (Magesh et al., 2025) to hiring in the labor market (Linnainmaa et al., 2021), and making earning forecasts (Campbell et al., 2026). In each of these settings, people must decide whether to rely on an algorithm or a human, and observed reliance may reflect either beliefs about accuracy or a preference for the source itself. Finance is a natural setting to study this question because advisor performance can be measured against objective outcomes.

Observed preferences for robo-advisors may arise from at least two distinct channels. First, individuals may hold biased beliefs about the relative competence of robo-advisors and human advisors. For example, they may believe that robo-advisors are more accurate, objective, or sophisticated, even when their actual performance is no better than that of human advisors. We refer to this channel as *belief bias*. Second, individuals may have a preference for advice from a particular source that is independent of its performance. That is, they may prefer a robo-advisor or a human advisor even after learning that the two sources provide advice of similar quality. We refer to this channel as *source preference*.

Distinguishing between these two channels is important because they have different implications for the adoption of robo-advice, the design of disclosure policies, and the broader use of automated decision aids. This distinction matters for firms and poli-

[1] See `https://adviserinfo.sec.gov/firm/summary/106715` for its latest Form ADV filed.

cymakers because correcting inaccurate beliefs requires information disclosure, whereas changing source preference may require different interventions, such as changes in design, communication, accountability, or the framing of advice.

Our paper makes three contributions. First, we develop and implement a dynamic design that decomposes demand for advice into belief bias and source preference, a distinction that field data and static experiments cannot make. Second, we show that static measures of preference for robo-advice conflate these channels: among subjects who initially preferred the robo-advisor, 31.88% in the *Odd-Even* treatment and 35.29% in the *Up-Down* treatment no longer exhibit this preference once accuracy is disclosed, yet a substantial subset retains it. Third, we show that, in the *Odd-Even* task (i.e., under risk), disclosure corrects belief bias but not source preference. On the other hand, in the *Up-Down* task (i.e., under uncertainty), disclosure of accuracy information reduces source preference for robo-advisor.

Identifying these two channels is difficult in naturally occurring settings. In practice, robo-advisors and human advisors may differ in ability, cost, transparency, consistency, reputation, personalization, and other attributes. As a result, observed reliance on one type of advisor over another does not reveal whether individuals prefer that advisor because they believe it performs better or because they have a taste for the source itself. Existing research has documented both affinity for and aversion to algorithms and robo-advisors, often emphasizing capabilities, trust, or decision quality as explanations for delegation decisions (e.g., D'Acunto et al., 2019; Brenner and Meyll, 2020; Back et al., 2023; Holzmeister et al., 2023). However, these explanations do not cleanly separate biased beliefs about advisor competence from source-based tastes.

We address this identification problem using a dynamic source preference framework. Source preference refers to the tendency of individuals to prefer one source of uncertainty over another, even when both sources are identical in terms of their probability distributions (Fox and Tversky, 1995; Chew and Sagi, 2008; Chew et al., 2008; Abdellaoui et al., 2011; Li, 2011; Chew et al., 2017; Li et al., 2018; Chen and Zhong, 2024). The source preference literature has typically studied static preferences over sources of uncertainty, such as familiar versus unfamiliar events or different ambiguity-generating processes. We extend this framework to the domain of advice and introduce a dynamic belief–taste

decomposition. The key idea is that static measures of source preference may confound tastes with biased beliefs. By eliciting valuations before and after subjects receive information about advisors' actual performance, our design allows us to distinguish belief bias from source preference.

Our experiment is designed to transform uncertainty about differences in advisor ability into objective risky events for which advisor accuracy is observable. In a pre-experiment, we collected predictions from both real-life investors, who serve as human advisors, and ChatGPT, which serves as the robo-advisor. These predictions were made in two treatments: *Odd-Even* and *Up-Down*. In the *Odd-Even* treatment, advisors predicted whether the last digit of a stock's closing price on a particular day was odd or even. In the *Up-Down* treatment, they predicted whether the stock's closing price moved up or down relative to the previous trading day.[2] Because all predictions concerned events that had already been realized, we could calculate the average prediction accuracy of each advisor type and use this information in the main experiment.

In the main experiment, subjects made a series of betting choices involving advice from the robo-advisor and the human advisor. The experiment consisted of two stages. In Stage 1, subjects made initial valuation decisions before receiving information about advisors' accuracy. In each round, they chose between taking a bet based on an advisor's advice and receiving a sure amount, using a menu-of-choice elicitation method similar to Abdellaoui et al. (2011). In Stage 2, subjects were informed about the actual accuracy of both advisors and were then given the opportunity to revise their initial choices. This pre/post disclosure structure is central to our identification strategy. Changes in valuations after accuracy disclosure identify the role of belief bias, while persistent valuation differences after disclosure identify source preference.

Our main findings reveal a novel pattern of changes in preferences. Before accuracy information is disclosed in the *Odd-Even* treatment, subjects are significantly more likely to prefer the robo-advisor over the human advisor, even though the underlying investments are objectively identical. This initial preference suggests that many subjects hold favorable beliefs about the robo-advisor's competence in the absence of performance information. However, once subjects learn that the robo-advisor and human advisor performed similarly, the proportion preferring the robo-advisor declines, while the pro-

[2]For each treatment, we randomly selected one real stock from U.S. stock exchanges.

portion preferring the human advisor increases. This shift indicates that part of the initial preference for the robo-advisor is driven by belief bias.

At the same time, the disclosure of accuracy information does not eliminate the preference for robo-advice. A substantial proportion of subjects continue to assign a premium to the robo-advisor even after learning that the two advisors performed similarly. This persistent preference cannot be explained by beliefs about relative performance and therefore identifies source preference: a preference for the robo-advisor itself, independent of its competence. Overall, the results show that both belief bias and source preference contribute to preferences for robo-advisors. We observe different patterns on the source preference in both *Odd-Even* and *Up-Down* treatments.

One may wonder whether providing accuracy information could also change source preference; in other words, whether source preference is unstable. This is indeed possible. To test it, we estimate the stability of the source function in Section 4.3. We find that source preferences are stable in the *Odd-Even* treatment but not in the *Up-Down* treatment.

These results suggest that source preference is stable when accuracy information carries little new content, as in the *Odd-Even* treatment where advice is correct about half the time, but becomes unstable when accuracy is ambiguous, as in the *Up-Down* treatment, where disclosure allows decision makers to update their beliefs about the source's competence.

In the literature, preference for a source of uncertainty (human vs. algorithm) is often treated as a stable attitude, such as algorithm aversion or appreciation. Our finding suggests that this stability depends on the task.

In the *Odd-Even* treatment, the advice is essentially a coin flip, so disclosed accuracy tells people little they did not already expect. Preferences between, say, a robo-advisor and a human advisor stay put because they reflect something intrinsic to the source (trust, comfort, algorithm aversion or appreciation) rather than beliefs about performance. On the other hand, in the *Up-Down* treatment (accuracy unknown): The advisor's quality is ambiguous, so disclosed accuracy is new information. People update their beliefs about how good each source is, and their preferences shift. Here, preference is partly an inference about competence, which fits the competence hypothesis of Heath and Tversky (1991)

that people's preference under uncertainty depends on perceived competence.

Related Literature. We contribute to the literature on robo-advice and algorithmic delegation by identifying two distinct mechanisms underlying preferences for robo-advisors. Existing studies have examined trust in algorithms, decision quality, familiarity, and reliance on automated advice (e.g., Hodge et al., 2021; Holzmeister et al., 2023; Tse et al., 2024). Our study differs by designing an environment in which advisor ability is controlled and by experimentally separating belief bias from source preference. To the best of our knowledge, this is the first study to propose and identify source preference as an explanation for preferences for robo-advisors over human advisors.

We contribute to the source preference literature by extending it from a static framework to a dynamic one. Prior work has documented source preference in settings such as familiar versus unfamiliar events, ambiguity, and effort provision (Fox and Tversky, 1995; Abdellaoui et al., 2011; Li et al., 2018; Chen and Zhong, 2024). In this literature, observed source preference is often attributed to tastes, under the assumption that individuals hold almost-objective probability beliefs (Machina, 2004). However, whether subjects actually hold unbiased beliefs is an empirical question. The theoretical framework of source preference allows both tastes and subjective beliefs to shape valuations (Chew and Sagi, 2008; Baillon et al., 2025). By observing how valuations change after performance information is disclosed, our design provides a way to decompose these channels empirically.

Our findings speak to research on investor beliefs, financial advice, and disclosure. Recent studies document that investors exhibit biases such as overconfidence, confirmation bias, and moral stock preferences (Statman et al., 2006; Pouget et al., 2017; Chew and Li, 2021; He and Li, 2026). Our results identify a different form of belief bias: biased beliefs about the relative competence of advice sources. In addition, our accuracy disclosure intervention is related to the financial education literature, which studies how objective information can correct biased beliefs and improve financial decision-making (Fernandes et al., 2014; Ambuehl et al., 2025). We show that disclosure can correct belief bias, but that it does not necessarily eliminate source-based tastes. This distinction is important for understanding when information provision can change delegation decisions and when preferences for advice sources may persist despite performance feedback.

Finally, our paper relates to a growing literature on the design of AI systems and on how individuals respond to AI as an advisor, evaluator, or overseer (e.g., Almog et al., 2024, 2026; Autor et al., 2025; Caplin et al., 2026; Choi et al., 2026; Kovach et al., 2026). This work studies how algorithms should be trained, and how people update beliefs, adjust effort, or change decisions when an algorithm is involved in the decision process. We instead hold advice quality fixed and ask why individuals prefer an algorithmic source in the first place, showing that part of this preference reflects a taste for the source rather than beliefs about its competence.

The remainder of the paper is organized as follows. Section 2 presents the theoretical background. Section 3 describes the experimental design. Section 4 reports the experimental results. Section 5 concludes.

2 Theoretical Background

Following the literature on source preference, we model each source as an algebra. There are two sources: H (Human Advisor) and R (Robo-Advisor). Throughout the paper, we assume each source is uniform so that subjects are probabilistically sophisticated with respect to each source (Savage, 1954; Abdellaoui et al., 2011).

To fix ideas, we consider the following simple setup. In our baseline treatment, the human advisor or the robo-advisor makes a prediction about the last digit of the stock price, and the subject, without knowing what prediction they make, needs to decide on their willingness to pay by following the advice from the human or the robot, respectively. With this setup, we can formulate a state space with an atomic event as a three-element vector. We let $\omega = (h, r, \ell)$ denote a generic atomic event, where $h \in \{o, e\}$ (odd or even) represents the human's prediction of the last digit, $r \in \{o, e\}$ represents the robot's prediction of the last digit, and $\ell \in \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$ denotes the realized last digit. Therefore, the state space Ω can be described as $\Omega := \{(h, r, \ell) \in \{o, e\} \times \{o, e\} \times \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$. Then, we formulate the human advisor as a source by the following algebra:

$$\mathcal{H} := \{\emptyset, E_H, \bar{E}_H, \Omega\}$$

where $\emptyset$ is the null event, E_H is the event that the human advisor makes a correct

prediction, i.e.,

$$E_H := \{(h, r, \ell) \in \Omega : h = e, \ell = \{0, 2, 4, 6, 8\}\} \cup \{(h, r, \ell) \in \Omega : h = o, \ell = \{1, 3, 5, 7, 9\}\},$$

and $\bar{E}_H$ is the complement of E_H. Here, E_H captures all the atomic events where the prediction h by the human matches the realization of the digit ℓ, regardless of the prediction by the robot, i.e., r. Note that probabilistic sophistication implies that there is a (subjective) probability measure P_H on the events over the algebra $\mathcal{H}$. Analogously, we model the robo-advisor as a source by an algebra, $\mathcal{R} := \{\emptyset, E_R, \bar{E}_R, \Omega\}$ where $E_R := \{(h, r, \ell) \in \Omega : r = e, \ell = \{0, 2, 4, 6, 8\}\} \cup \{(h, r, \ell) \in \Omega : r = o, \ell = \{1, 3, 5, 7, 9\}\}$, and $\bar{E}_R$ is the complement of E_R. Hence, by probabilistic sophistication, we assume the existence of the subjective probability measure P_R over the space $\mathcal{R}$.

Let $(E^1 : x_1, \ldots, E^n : x_n)$ denote a prospect yielding outcome x_c if event E^c happens for $c = 1, \ldots, n$. In the paper, we are primarily interested in two-outcome prospects of the two following forms, for $i = H, R$ and for some $x, y \in \mathbb{R}_+$ with $x > y$

$$(E_i : x, \bar{E}_i : y).$$

Here, the prospect represents the bet on the source $i = H, R$, where the subject earns an amount x if the advisor i makes a correct prediction, and earns y otherwise. For simplicity, we denote each binary prospect by xE_iy, where the outcome is x if event E_i occurs, the outcome is y otherwise. In source theory, probabilistic sophistication implies that there exists a weighting function $w_i : [0, 1] \to [0, 1]$ for $i = H, R$ so that each prospect can be mapped to an evaluation, where the weighting function represents each agent's source preference. Following the literature, we call w_i the source function.

Therefore, for each source $i = H, R$, we have

$$xE_iy \mapsto w_i(P_i(E_i))u(x) + (1 - w_i(P_i(E_i)))u(y) \tag{1}$$

where $u : \mathbb{R} \to \mathbb{R}$ represents the Bernoulli utility of the agent.

The novelty of our experimental design is that we can resolve the uncertainty regarding the true underlying probability of each advisor making a correct prediction. As a result, the DM has the knowledge of the objective probability measures over the algebra $\mathcal{H}$ and $\mathcal{R}$. This transforms the subject's decision problem into a standard decision-making

problem under risk, where the probability is known. This allows us to disentangle the effect of potential belief bias and source preference. For $i = H, R$, let $P(E_i)$ be the objective probability of the event E_i. Therefore, after revising their belief and assuming they follow the objective probability, each prospect is evaluated according to

$$xE_iy \mapsto w_i'(P(E_i))u(x) + (1 - w_i'(P(E_i)))u(y) \tag{2}$$

where the source function w_i' is applied over the objective probability.

With these formalizations, we state the following properties of interest. First, we introduce the definition of belief bias. Since our setup effectively delineates the (prior) subjective probability and the objective probability, belief bias is defined from the primitives of the model.

Definition 1 (Belief Bias)**.** A subject is said to exhibit a positive belief bias for source i if $P_i(E_i) > P(E_i)$. A subject is said to exhibit a negative belief bias for source i if $P_i(E_i) < P(E_i)$.

Secondly, we would like to perform inter-source comparison over the belief bias and the source preferences. First, we define the following percentage change as a basis of measure. In particular, we denote B_i as the relative degree of belief bias by comparing the subjective belief's $P_i(E_i)$ and the objective probability $P(E_i)$. Therefore, we have

$$B_i := \frac{P_i(E_i) - P(E_i)}{P(E_i)}$$

We define our comparative statics with respect to this measure.

Definition 2 (Comparative Belief Bias)**.** Let $i, j \in \{H, R\}$ and $i \neq j$. A subject is said to have a favorable belief bias for source i over source j if $B_i > B_j$.

Moreover, we are interested in comparing source preference. Source preference is captured by the source function w_i, w_i' for $i = H, R$. Therefore, to examine whether a subject's preference for each source can be attributed to the source function itself, we analyze w_i directly. To this end, following the literature, we consider a specification of the source function. In particular, we assume that the source function $w_i, w_i' : \mathbb{R} \to \mathbb{R}$ will embed the parameters in the following forms.

$$W(x; a, b) = (\exp(-(-\ln(x))^a))^b$$

Hence, we have $w_i(x) := W(x; \alpha_i, \beta_i)$ and $w_i(x) := W(x; \alpha'_i, \beta'_i)$. Here, for $i = H, R$, α_i and α'_i represent, respectively, the level of likelihood insensitivity before information disclosure and after information disclosure; β_i and β'_i represents, respectively, the degree of pessimism before information disclosure and after information disclosure. Previous literature (Abdellaoui et al., 2011) has estimated these two parameters for uncertain and risky sources. In our setting, similarly, both sources are uncertain before information disclosure and become risky after information disclosure. For ease of comparison, we focus on the pessimism parameter that differentiates their attitudes toward the sources in terms of how they "underweight" the outcome. Therefore, we assume that $\alpha_H = \alpha_R = \alpha'_H = \alpha'_R$ for each subject and estimate β_i and β'_i for $i = H, R$. Hence, it gives us a parsimonious model for defining the source preference for subjects, so that we can compare the source function based on one parameter difference across the two sources.

Definition 3 (Comparative Source Preference)**.** Let $i, j \in \{H, R\}$ and $i \neq j$. The subject is said to exhibit a source preference under uncertainty for i over j if $\beta_i < \beta_j$. The subject is said to exhibit a source preference under risk for i over j if $\beta'_i < \beta'_j$.

Based on our specification, for each individual, we will need to estimate $\beta_i, \beta'_i, \alpha, P_i(E_i)$ for $i = H, R$. We relegate the estimation procedure to Section A.1 in the Appendix.

2.1 Hypotheses

Before the provision of accuracy information, there is no reason to expect the robo-advisor and the human advisor to differ in accuracy. Accordingly, expected utility theory predicts that decision makers should be indifferent between the two types of advisors. This leads to the following hypothesis:

Hypothesis 1. Prior to the provision of accuracy information, subjects will be indifferent between the robo-advisor and the human advisor.

Note that an alternative hypothesis is that some subjects may exhibit a source preference. If some subjects exhibit a strict preference at this stage, this may be attributed either to belief bias or to a source preference. The accuracy information provided later will allow us to distinguish between these two possibilities.

According to expected utility theory, decision makers should prefer the advisor with higher objective accuracy. In our experiment, the announced accuracy of the robo-advisor was lower than that of the human advisor in both treatments. Expected utility theory

therefore predicts that, once accuracy information is disclosed, decision makers should prefer the human advisor.

Consequently, a preference for the robo-advisor after accuracy is revealed can be attributed to a source preference. In contrast, if a decision maker shifts from preferring the robo-advisor to preferring the human advisor (or to being indifferent) once accuracy is revealed, this would suggest that the initial preference was driven by belief bias.

This reasoning yields the following hypothesis, grounded in expected utility theory:

Hypothesis 2. After information about advisor accuracy is disclosed, subjects will strictly prefer the advisor with higher accuracy.

Note that an alternative hypothesis is that some subjects may continue to prefer the robo-advisor due to a source preference.

Finally, since we implement the *Odd-Even* and *Up-Down* tasks as separate treatments, we expect source preferences to behave differently across these two settings. In particular, since it is reasonable to believe that the *Odd-Even* task is relatively more objective, we therefore expect the source preference parameter to remain relatively stable following the disclosure of accuracy information. Accordingly, for a given individual, the estimated β_i and β_i' should not differ much.

Hypothesis 3. In the *Odd-Even* treatment, the source function is stable before and after information disclosure.

The testable implication of this hypothesis in our specification is a statistical test of the difference in the estimated β_i and β_i'.

3 Experimental Design

We conducted two experiments: a pre-experiment and a main experiment, each of which included two treatments: *Odd-Even* and *Up-Down*. These treatments respectively involved predictions or advice on whether the last digit of a real stock's closing price was odd or even (*Odd-Even*), and whether the stock's price closed up or down compared to the previous trading day (*Up-Down*).

As shown in Figure 1, the pre-experiment was designed to evaluate the accuracy of investment predictions made by both human advisors (played by real-life investors) and

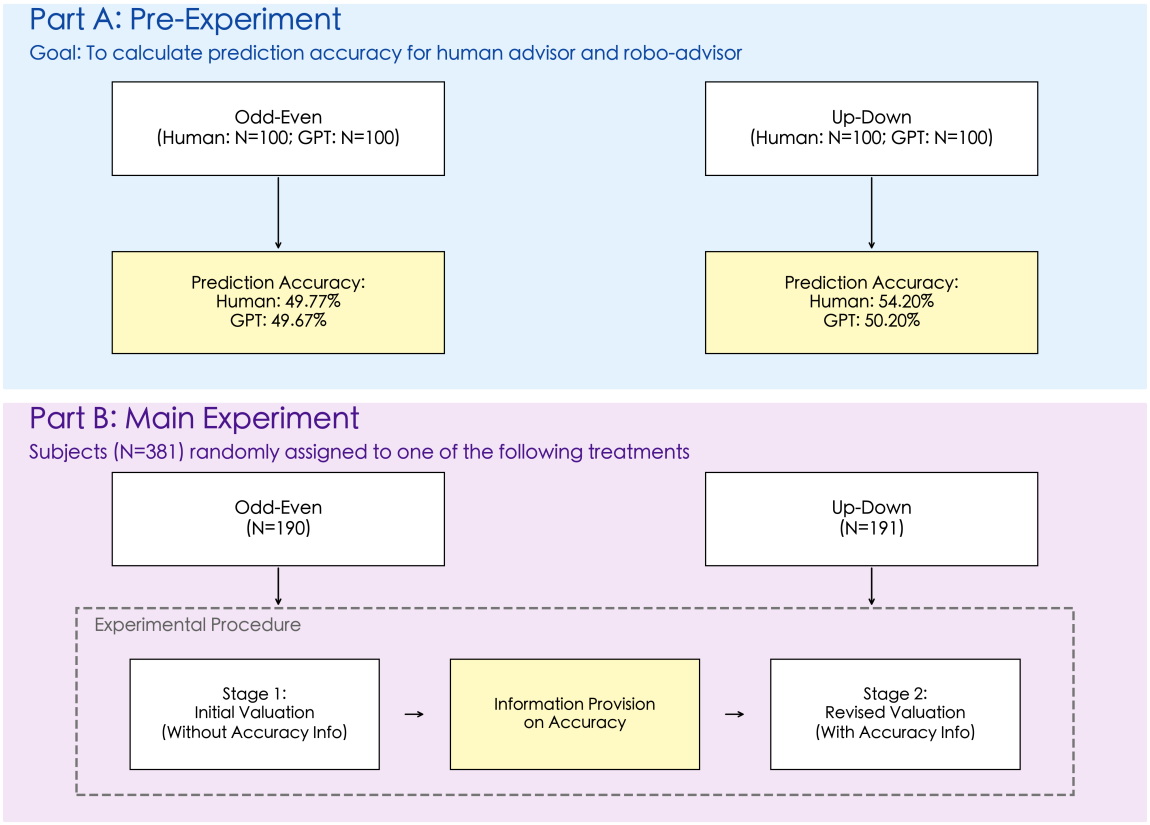

Figure 1: Experimental Design

robo-advisors (played by ChatGPT).[3] The results on prediction accuracy were subsequently used in the main experiment, which comprised two key tasks: an initial valuation task and a revised valuation task. These tasks were implemented to elicit subjects' valuations (measured by the certainty equivalent) of either type of advisor's advice, both before and after they were informed of its prediction accuracy. Both experiments were programmed in oTree (Chen et al., 2016).[4]

3.1 Pre-Experiment

The objective of the pre-experiment was to collect investment predictions from both real-life investor subjects (human advisors) and ChatGPT (robo-advisors). These predictions were then used to calculate each advisor's average prediction accuracy, which was subsequently presented to subjects in the revised valuation task of the main experiment.

[3]The pre-experiment was conducted on July 15, 2025, using OpenAI's o3-pro model for the robo-advisors.

[4]We provide the full set of instructions in the Online Appendix (Sections B and C).

3.1.1 Procedure

We divided the pre-experiment into two sections: one for human advisors (human advisor section) and the other for robo-advisors (robo-advisor section). In the human advisor section, we recruited 200 real-life American investors via Prolific.[5][6]

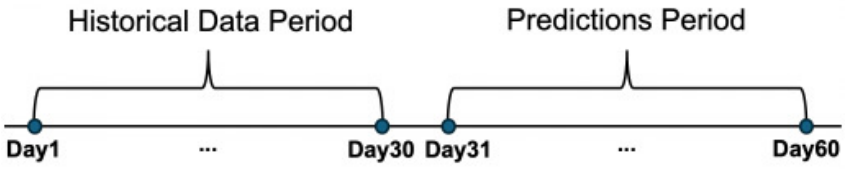

Figure 2: Timeline Structure for the Prediction Task

Human Advisor Section. The main task involved predicting the performance of a real stock listed on the stock exchange. At the beginning of the task, subjects first received historical data on a stock's closing prices (or closing price changes compared to the previous trading day in the *Up-Down* treatment) for the first 30 consecutive trading days (labelled "Day 1" to "Day 30", see Figure 2) from an undisclosed year. Based on this information, they were asked to make predictions: odd-even, up-down, or the exact last digit of the closing price, depending on the assigned treatment, for the subsequent 30 trading days ("Day 31" to "Day 60") of the same year. An example from the *Odd–Even* treatment is illustrated in Figure 3.

The stock was randomly selected from the exchange, remained the same throughout the experiment, and all uncertainty about its realized outcomes had already been resolved at the time of subjects' participation.[7] For each prediction, subjects selected one option (odd or even in the *Odd-Even* treatment; up or down in the *Up-Down* treatment). After the experiment, one prediction from Day 31 to Day 60 was randomly drawn for bonus payment. If the prediction for the drawn day was correct, the subject earned 2 USD; otherwise, the payment was 0.

[5]Subjects were required to use a desktop computer to complete the study. We screened subjects for U.S. nationality and prior investment experience, and we ensured a balanced gender ratio in the sample. Subjects were randomly and evenly assigned to the two treatments (100 per treatment). Each subject made independent predictions. In the robo-advisor section, we provided the same instructions to ChatGPT and enabled it to independently generate 30 predictions on 100 separate occasions.

[6]We used the "Temporary Chat" function to achieve this.

[7]Although the stocks were randomly chosen, the historical data was deliberately screened. Specifically, we selected stocks with exactly 15 odd and 15 even last digits over 30 consecutive trading days for the *Odd-Even* treatment; 15 up and 15 down days for the *Up-Down* treatment.

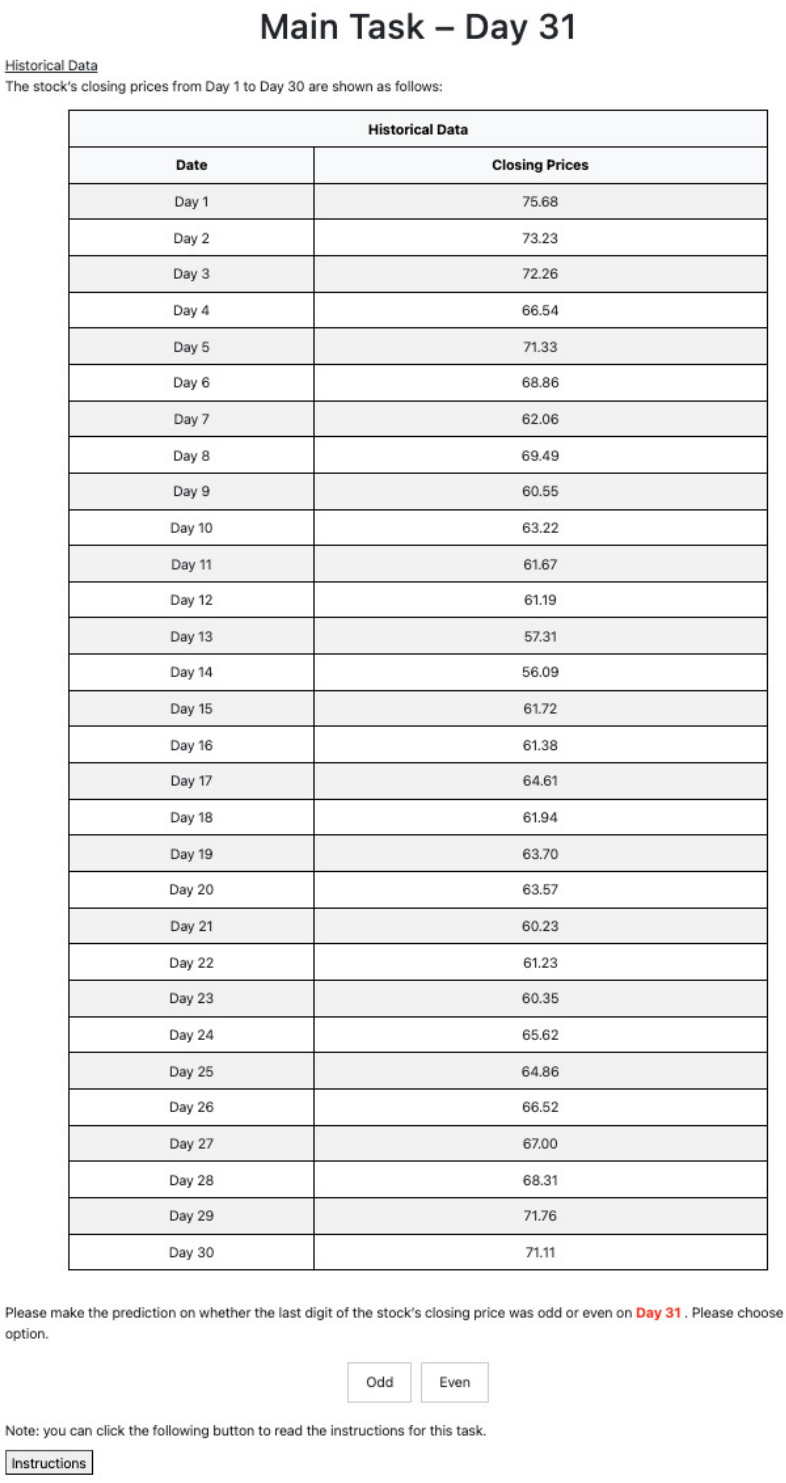

Main Task – Day 31

Historical Data

The stock's closing prices from Day 1 to Day 30 are shown as follows:

Historical Data	
Date	**Closing Prices**
Day 1	75.68
Day 2	73.23
Day 3	72.26
Day 4	66.54
Day 5	71.33
Day 6	68.86
Day 7	62.06
Day 8	69.49
Day 9	60.55
Day 10	63.22
Day 11	61.67
Day 12	61.19
Day 13	57.31
Day 14	56.09
Day 15	61.72
Day 16	61.38
Day 17	64.61
Day 18	61.94
Day 19	63.70
Day 20	63.57
Day 21	60.23
Day 22	61.23
Day 23	60.35
Day 24	65.62
Day 25	64.86
Day 26	66.52
Day 27	67.00
Day 28	68.31
Day 29	71.76
Day 30	71.11

Please make the prediction on whether the last digit of the stock's closing price was odd or even on **Day 31** . Please choose one option.

Odd Even

Note: you can click the following button to read the instructions for this task.

Instructions

Figure 3: Main Task in the Pre-Experiment

To ensure comprehension, all subjects completed two practice tasks before proceeding to the main tasks. To mitigate order effects, the options (in *Odd-Even* and *Up-Down*) were presented in random order. During the prediction tasks, subjects were allowed to review the instructions at any time. After the prediction tasks, subjects completed a questionnaire collecting their demographics, including gender, age, income, education, and employment status.

Robo-Advisor Section. As detailed above, we provided ChatGPT with the same instructions shown to human subjects and enabled it to independently generate 30 predictions on 100 separate occasions per treatment. This yielded a total of 3,000 predictions, matching the number obtained from human advisors. By combining the two datasets, we calculated the average prediction accuracy for each type of advisor, conditional on each treatment, which was then used in the main experiment.

3.2 Main Experiment

We conducted the main experiment to elicit subjects' valuations of advice from both the robo-advisor and the human advisor, measured before and after disclosing the prediction accuracy from the pre-experiment. The treatments in the main experiment correspond directly to those in the pre-experiment. Each subject completed two stages of decision-making (Stage 1: Initial Valuation Task; Stage 2: Revised Valuation Task), with each stage consisting of five decision rounds. Throughout the experiment, all amounts were displayed in Experimental Currency Units (ECU). At the end of the experiment, bonuses were converted into USD at a rate of 1 USD = 10 ECU. Subjects' final payoffs consisted of a participation fee plus the bonus from a randomly selected task.

3.2.1 Procedure

Stage 1: Initial Valuation Task. In the initial valuation task, taking the *Odd–Even* treatment as an example, subjects made a series of choices. As shown in Figure 4, subjects made decisions using a multiple price list between following an advisor's advice (Option 1) and receiving a sure amount (Option 2). Specifically, in each item, subjects decided whether to follow the advice of an advisor (either a robo-advisor or a human advisor) regarding whether the last digit of a stock's closing price on a randomly drawn trading day (between Day 31 and Day 60) was odd or even, or instead choose the safe option. Before

making decisions, subjects were informed about how advisors generated their predictions:

- **Robo-Advisor:** We provided ChatGPT with historical data on the stock's closing prices for the trading days from Day 1 to Day 30, without disclosing the stock's ticker symbol. The ChatGPT provided 100 independent predictions on whether the last digit (i.e., the hundredths decimal place) of the stock's closing price was odd or even for each trading day from Day 31 to Day 60. For each prediction, the ChatGPT will choose one option (either odd or even) and if its choice matches the actual digit's parity (i.e., odd or even), the prediction is considered correct.

- **Human Advisor:** We provided 100 real-life investors with historical data on the stock's closing prices for the trading days from Day 1 to Day 30, without disclosing the stock's ticker symbol. Each investor independently provided predictions on whether the last digit (i.e., the hundredths decimal place) of the stock's closing price was odd or even for each trading day from Day 31 to Day 60. For each prediction, the investor will choose one option (either odd or even) and if her/his choice matches the actual digit's parity (i.e., odd or even), the prediction is considered correct.

Subjects completed five rounds of decisions. In each round, a trading day between Day 31 and Day 60 was randomly drawn to determine which prediction would be used for bonus payment calculations. In each round, subjects' valuations of advice were elicited using a menu-of-choice design. Specifically, subjects chose between the following options:

- **Follow an Advisor's Advice.** If a subject chooses to follow an advisor, one of the 100 predictions from the advisor on the drawn date will be randomly drawn as her/his prediction. If the prediction is correct for the randomly drawn trading day from Day 31 to Day 60, the subject will win the specified amount listed on the decision page; otherwise, the subject will receive the lower amount listed on the same page.

  We introduced five different payoff pairs for the bets (i.e., lotteries): (10, 0), (15, 0), (25, 0), (20, 10), and (25, 15). To mitigate potential order effects, both the lotteries and the blocks eliciting subjects' valuations of advice from the robo-advisor and the human advisor were presented in random order. The higher value indicates

the payoff a subject received if she followed the chosen advisor's advice and the prediction was correct; the lower value indicates the payoff otherwise.

- **Receive a Sure Amount.** The subject will receive a sure amount listed on the decision page (see an example of lottery (10, 0) in the Figure 4).

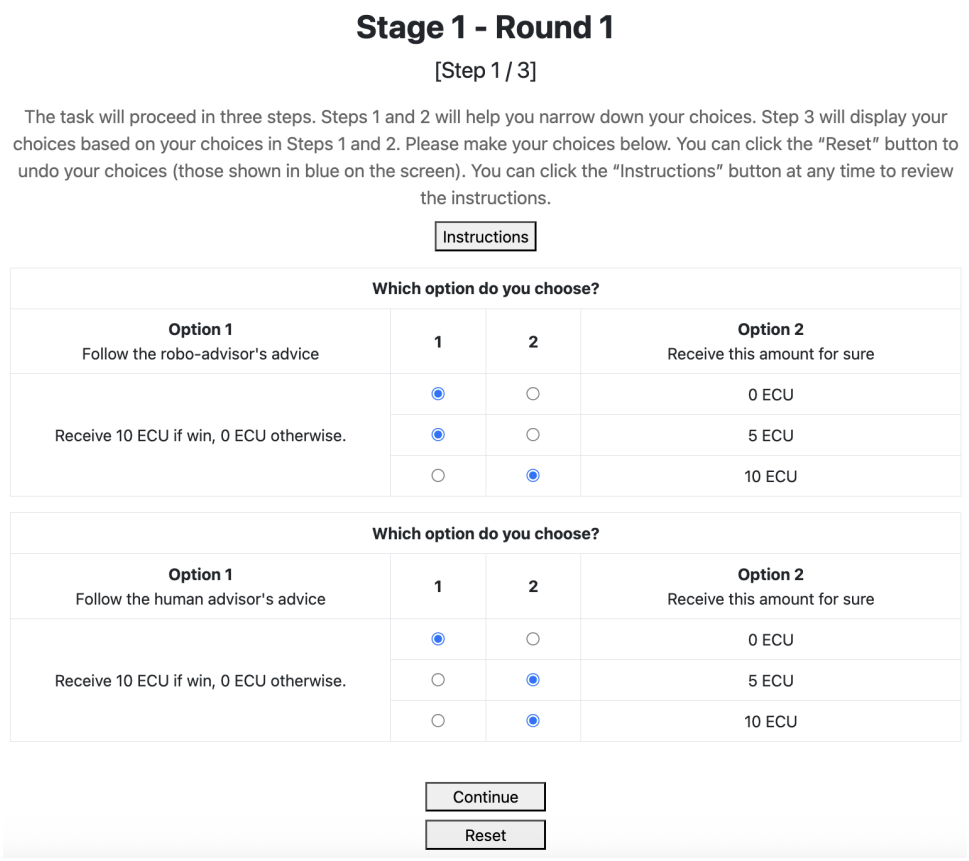

Figure 4: Initial Valuation Task in the Main Experiment

As shown in Figure 4, each round in Stage 1 proceeded in three steps. Steps 1 and 2 guided subjects to narrow down their choices, while Step 3 displayed their final options based on these preliminary choices. Subjects could return to any previous step and click the "Reset" button to undo their selections and review the instructions at any time before submitting a final decision in Step 3.

Stage 2: Revised Valuation Task. In Stage 2, subjects again completed five rounds of decisions. In each round, the subject's Stage 1 choice list was displayed alongside information about advisors' prediction accuracy (see Figures 1 and 5). Subjects were then allowed to revise their Stage 1 decisions. Subjects were informed that bonus payments would be determined based on Stage 2 decisions.



Stage 2 - Round 1

We would like to provide you with information about the advisors' prediction accuracy:

Based on the data, on average, **49.67%** of the **robo-advisor**'s predictions are correct on any given trading day.

Based on the data, on average, **49.77%** of the **human advisor**'s predictions are correct on any given trading day.

Your previous decisions in Stage 1 on whether to follow an advisor's advice or receive a sure amount are displayed on this page. You can either update your decisions or keep the original ones. Once you have finished your revisions or decided to keep your previous decisions, please click the "Submit" button. If you want to undo your changes and restore your original choices, you may click the "Reset to Previous Decisions" button.

Note: you can click the following button to read the instructions for this task.

Instructions

Which option do you choose?

Option 1 Follow the human advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 10 ECU if win, 0 ECU otherwise.	◉	○	0 ECU
	○	◉	0.5 ECU
	○	◉	1 ECU
	○	◉	1.5 ECU
	○	◉	2 ECU
	○	◉	2.5 ECU
	○	◉	3 ECU
	○	◉	3.5 ECU
	○	◉	4 ECU
	○	◉	4.5 ECU
	○	◉	5 ECU
	○	◉	5.5 ECU
	○	◉	6 ECU
	○	◉	6.5 ECU
	○	◉	7 ECU
	○	◉	7.5 ECU
	○	◉	8 ECU
	○	◉	8.5 ECU
	○	◉	9 ECU
	○	◉	9.5 ECU
	○	◉	10 ECU

Figure 5: Revised Valuation Task in the Main Experiment

A total of 393 real-life American investors recruited via Prolific participated in our main experiment. Each subject was randomly assigned to one of two treatments and participated in only one of them. Subjects who did not submit a complete choice list in the revised valuation task (12 subjects in total) were dropped from the sample. Thus, in our sample for analysis, there are 190 subjects in the *Odd-Even* treatment and 191 in the *Up-Down* treatment. Each subject received a flat payment of 2 USD, along with a



Figure 5: Revised Valuation Task in the Main Experiment (Continued)

potential payoff from a randomly drawn task based on their individual decisions. Prior to the experiment, all subjects were required to give consent and were informed that their decisions and personal information would be treated with strict confidentiality and kept anonymous throughout the study.

Knowledge of FinTech. We measured subjects' knowledge of FinTech using a 5-point Likert scale, comprising five items on mainstream financial technologies: Artificial Intelligence, Blockchain, Cloud Computing, Big Data, and Digital Currency (*Cronbach's* $\alpha = 0.84$). Subjects rated their familiarity with each technology on a scale of 1 (very unfamiliar) to 5 (very familiar). We computed the sum of responses to all five items as a measure of each subject's knowledge of FinTech.

Familiarity and Prior Experience. Subjects rated their familiarity with both types of advisors on a scale of 1 (very unfamiliar) to 10 (very familiar). Additionally, they were asked about any prior experience with robo-advisors and human advisors. If a subject reported prior experience, she was then asked to rate her feelings about that experience on a scale of 1 (very bad) to 10 (very good), reflecting the quality of her experiences with that type of advisor. In addition to the two tasks mentioned above, subjects in the main experiment also completed a questionnaire collecting their demographic information, including gender, age, income, education, and employment status.

4 Experimental Results

In this section, we report the main results based on subjects' valuations elicited both before and after the information update. In particular, we examine subjects' preferences between robot and human advice before the information update in Section 4.1, as well as their preferences for a specific advisor across the initial and revision tasks in Section 4.2. Finally, we report evidence from the parametric estimations to further examine the role of belief bias and source preference in Section 4.3.

4.1 Preference for Advisor before Information Update

We first examine subjects' preferences between the robo-advisor and the human advisor before objective accuracy information is disclosed. At this stage, subjects do not know the advisors' objective probabilities of success. Their valuations may therefore reflect both their beliefs about advisor accuracy and any intrinsic preference over the source of advice. We refer to this setting as one involving uncertainty about advisor accuracy.

Table 1: Valuations for Robo vs. Human Advisor before Information Update

	Mean Valuation		Difference	p-value	
	Robot	Human		t-test	Wilcoxon
Odd-Even	10.47	10.29	0.18	0.20	0.01***
Up-Down	10.41	10.33	0.08	0.55	0.04**

Notes: The p-values are based on paired t-tests and Wilcoxon signed-rank tests. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Valuations across Advisors. Table 1 compares subjects' valuations of the robo-

advisor and the human advisor in both treatments. In the *Odd-Even* treatment, the mean valuation of the robo-advisor is 10.47, compared with 10.29 for the human advisor. This difference is not statistically significant ($p = 0.20$, paired t-test), while the distributions of valuations differ significantly ($p < 0.01$, Wilcoxon signed-rank test).[8] A similar pattern is observed in the *Up-Down* treatment: the mean valuation is 10.41 for the robo-advisor and 10.33 for the human advisor, with no significant difference in means ($p = 0.55$), but a significant difference in their distributions ($p = 0.04$).

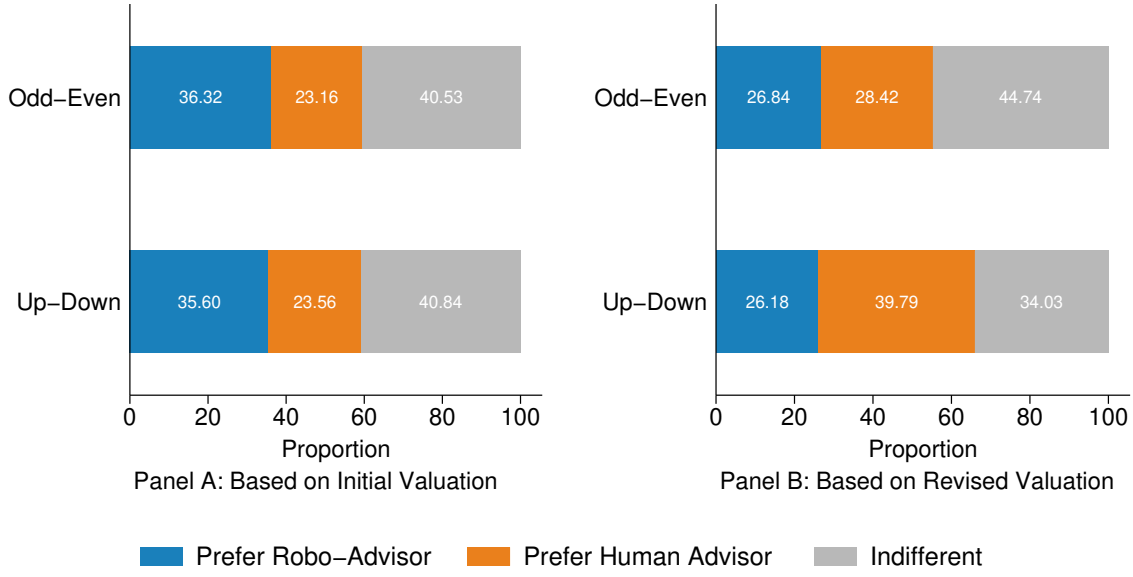

Figure 6: Decomposition of Preferences in Initial vs. Revision Tasks

Notes: The analysis sample includes 190 subjects in the *Odd-Even* treatment and 191 in the *Up-Down* treatment. Panel A reports the proportion of subjects who preferred the robo-advisor, preferred the human advisor, or were indifferent between the two based on their decisions in the initial valuation task. Panel B reports the corresponding proportions based on their decisions in the revised valuation task.

Thus, it implies that, in both treatments, valuations at the aggregate level mask substantial heterogeneity in subjects' relative valuations of the two advisors. We therefore turn to individual-level preferences. We classify a subject as preferring the robo-advisor if her mean valuation across the five valuation tasks is higher for the robo-advisor than for the human advisor, as preferring the human advisor if the reverse is true, and as indifferent otherwise. Figure 6 and Table 2 report the resulting classification.

[8]Unless otherwise indicated, statistical significance is evaluated at the 5% level. We use paired t-tests and Wilcoxon signed-rank tests for within-subject comparisons.

Table 2: Proportion Preferring Robo vs. Human Advisor before Information Update

	Proportion		Difference	p-value
	Robot	Human		
Odd-Even	36.32%	23.16%	13.16%	0.01***
Up-Down	35.60%	23.56%	12.04%	0.01***

Notes: The p-values are based on two-sample t-tests of proportions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Before the information update, subjects are significantly more likely to prefer the robo-advisor than the human advisor in both treatments. In the *Odd-Even* treatment, 36.32% of subjects prefer the robo-advisor, compared with 23.16% who prefer the human advisor, a difference of 13.16 percentage points ($p < 0.01$). The corresponding proportions in the *Up-Down* treatment are 35.60% and 23.56%, respectively, yielding a similarly significant difference of 12.04 percentage points ($p < 0.01$).

Taken together, the two treatments reveal a similar pre-disclosure pattern. Although mean valuations do not differ significantly between the two advisors, significantly more subjects favor the robo-advisor than the human advisor. We therefore reject Hypothesis 1 and find evidence of source-dependent preferences before objective accuracy is disclosed. At this stage, however, these comparisons cannot distinguish whether the relative preference for the robo-advisor is driven by biased beliefs about advisor accuracy or by a source preference.

Result 1. Before the information update, significantly more subjects prefer the robo-advisor over the human advisor in both the *Odd-Even* and *Up-Down* treatments.

4.2 Preference for Advisor after Information Update

We next examine how subjects respond when objective accuracy information is disclosed. The information update resolves uncertainty about the advisors' success probabilities by revealing the objective accuracies measured in the pre-experiment. As a result, the comparison before and after the information update allows us to identify the role of prior beliefs, while post-disclosure choices take place under known objective probabilities (i.e., under risk).

Changes in Valuations. Table 3 reports valuations before and after the information update. A common pattern emerges across the two treatments: valuations of the robo-

Table 3: Valuations for Advisor in Initial vs. Revision Task

	Mean Valuation		Difference	p-value	
	Initial	Revision		t-test	Wilcoxon
Robo-Advisor:					
Odd-Even	10.47	10.14	0.33	0.00***	0.00***
Up-Down	10.41	9.94	0.47	0.00***	0.00***
Human Advisor:					
Odd-Even	10.29	10.31	-0.02	0.86	0.31
Up-Down	10.33	10.38	-0.05	0.65	0.68

Notes: The p-values are based on paired t-tests and Wilcoxon signed-rank tests. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

advisor decline significantly after objective accuracy is revealed, whereas valuations of the human advisor remain approximately unchanged.

In the *Odd-Even* treatment, the mean valuation of the robo-advisor decreases from 10.47 to 10.14 ($p < 0.01$), while the mean valuation of the human advisor changes only from 10.29 to 10.31 and is not statistically different from its initial level. Similarly, in the *Up-Down* treatment, the mean valuation of the robo-advisor falls from 10.41 to 9.94 ($p < 0.01$), whereas the valuation of the human advisor remains stable, changing from 10.33 to 10.38.

The asymmetric response to information is consistent with favorable prior beliefs about the robo-advisor under the assumption of stability of source function. Before disclosure, subjects appear to value the robo-advisor partly on the basis of beliefs that are revised downward once its objective accuracy becomes known. In contrast, the absence of an aggregate change in valuations of the human advisor suggests substantially less systematic belief updating in one direction if the source function is stable. Note that this discussion is based on the stability of source function, which we are able to test at the population level in the next section.

Robo-Advisor Premium. To examine changes in relative preferences between the two advisors, we define the robo-advisor premium as the difference between a subject's valuation of the robo-advisor and that of the human advisor. A positive (negative) premium indicates a higher valuation of the robo-advisor (human advisor).

In the *Odd-Even* treatment, the mean robo-advisor premium decreases from 0.17 in the initial task to -0.18 in the revision task, a statistically significant decline of 0.35 ($p <$

0.01). A similar pattern is observed in the *Up-Down* treatment, where the mean premium decreases from 0.08 to −0.45, a significant decline of 0.53 ($p < 0.01$). Thus, in both treatments, the average robo-advisor premium reverses sign following the information update, indicating a shift in relative valuations from favoring the robo-advisor to favoring the human advisor.

At the individual level, among subjects with an initially positive robo-advisor premium, 31.88% in the *Odd-Even* treatment and 35.29% in the *Up-Down* treatment no longer exhibit this preference following the information update. However, a subset of subjects continues to value the robo-advisor more highly than the human advisor.

Result 2. The preference for human-advisor is stable under risk and uncertainty, while the preference for robo-advisor was significantly reduced when uncertainty about the performance was resolved.

Table 4: Proportion Preferring Advisor in Initial vs. Revision Task

	Proportion		Difference	p-value
	Initial	Revision		
Preferring Robo-Advisor:				
Odd-Even	36.32%	26.84%	9.48%	0.05**
Up-Down	35.60%	26.18%	9.42%	0.05**
Preferring Human Advisor:				
Odd-Even	23.16%	28.42%	-5.26%	0.24
Up-Down	23.56%	39.79%	-16.23%	0.00***

Notes: The p-values are based on two-sample t-tests of proportions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Changes in Preferences. The information update also produces a systematic shift in subjects' relative preferences across advisors. As shown in Figure 6 and Table 4, the proportion of subjects preferring the robo-advisor declines in both treatments following disclosure.

In the *Odd-Even* treatment, the share preferring the robo-advisor decreases from 36.32% before the information update to 26.84% afterward, a decline of 9.48 percentage points ($p < 0.05$). At the same time, the share preferring the human advisor increases from 23.16% to 28.42%, although this increase is not statistically significant. In the *Up-Down* treatment, the proportion preferring the robo-advisor similarly decreases from

35.60% to 26.18% ($p < 0.05$), while the proportion preferring the human advisor rises substantially from 23.56% to 39.79% ($p < 0.01$).

These results show that resolving uncertainty about advisor accuracy shifts preferences away from the robo-advisor and toward the human advisor in both treatments. The shift is particularly pronounced in the *Up-Down* treatment, where the increase in preference for the human advisor is substantially larger.

Preferences under Risk. The post-disclosure decisions provide a useful benchmark for distinguishing belief-driven preferences from source preference. After the information update, subjects know the objective success probabilities of both advisors. Moreover, the human advisor has higher objective accuracy than the robo-advisor in both treatments. Under expected utility theory and in the absence of source-dependent tastes, subjects should therefore weakly prefer the human advisor once these probabilities are known.

The decline in preference for the robo-advisor following disclosure is consistent with belief bias accounting for part of its initial advantage, and/or a change in source function after disclosure. However, the information update does not eliminate preference for the robo-advisor. Even after objective probabilities are revealed, 26.84% of subjects in the *Odd-Even* treatment and 26.18% in the *Up-Down* treatment continue to prefer the robo-advisor, thus rejecting hypothesis 2. These persistent preferences cannot be explained by favorable beliefs about its objective accuracy once that accuracy is known and therefore can be consistent with source preference.

Taken together, the valuation evidence suggests two distinct mechanisms. First, belief bias might contribute to the initial relative preference for the robo-advisor, as indicated by the significant downward revision in its valuation following the information update. Second, a non-negligible fraction of subjects continue to prefer the robo-advisor even under known objective probabilities, providing evidence of source preference. In the next subsection, we use our estimation to separate these mechanisms more directly and to examine how they differ across the *Odd-Even* and *Up-Down* treatments.

Result 3. Information disclosure shifts preferences away from the robo-advisor and toward the human advisor in both treatments. Nevertheless, a substantial fraction of subjects continue to prefer the robo-advisor after objective accuracies are revealed, consistent with the presence of source preference.

4.3 Belief, Taste and Source Preference

In this section, to further explore the mechanisms driving subjects' preferences for advisors, we perform an estimation of our model as discussed in Section A.1. Note that, due to the discrepancy in accuracy, we cannot attribute source preference to any subjects beyond those who still prefer the robo-advisor even after learning the accuracy information. Therefore, to get a better sense of the true underlying distribution of source preference for each advisor, it is necessary to employ the estimation to compare the underlying taste parameter in driving the source preference as discussed in Definition 3. Moreover, our estimation also allows us to look further into the subjective probabilities and estimate the level of belief bias across subjects. This section reports the results both from the *Odd-Even* and *Up-Down* treatments.

Source Function Stability. Following our theoretical framework, we first estimate the pre-disclosure shifts, δ_R and δ_H, in the pessimism parameter for the robot and human sources, respectively. In the *Odd-Even* treatment, neither estimated shift is statistically different from zero ($\delta_R = -0.04$, $p = 0.08$; $\delta_H = -0.01$, $p = 0.68$). These results are consistent with the stability of the source function, supporting hypothesis 3. In contrast, in the *Up-Down* treatment, the estimated shift for the robot source is negative and statistically significant ($\delta_R = -0.04$, $p = 0.03$), whereas the estimated shift for the human source is not statistically different from zero ($\delta_H = -0.03$, $p = 0.23$).

In the *Odd-Even* treatment, the advice is like a coin-flip (with known-probability) (thus can be considered as the case of decision making under risk), the disclosed accuracy information carries little new information. Hence, the source preference is stable and it reflects something intrinsic to the source (trust, comfort, algorithm aversion or appreciation) rather than beliefs about performance. However, in the *Up-Down* treatment, the advisor's quality is ambiguous, so disclosed accuracy is new information. The disclosure allows decision makers to update their beliefs about the source's competence, supporting the competence hypothesis of Heath and Tversky (1991). Thus, our result suggests that preference for robo-advisor depends on whether the decision-context is under risk or uncertainty.

Result 4. Source preference is stable in the *Odd-Even* task (i.e., under risk), while it is not stable in the *Up-Down* task (i.e., under uncertainty).

Belief Bias. Figure 7 illustrates the distributions of subjects' estimated prior subjective probabilities in the *Odd-Even* and *Up-Down* treatments, respectively. Table 5 summarizes the corresponding belief-bias measures for both treatments.

In the *Odd-Even* treatment, subjects held a significantly higher prior subjective probability for the robo-advisor than for the human advisor (mean $p_R^S = 50.59\%$ versus mean $p_H^S = 49.66\%$; $p < 0.01$, paired t-test). Relative to the corresponding objective probabilities, subjects significantly overestimated the accuracy of the robo-advisor ($p_R^S = 50.59\% > p_R^O = 49.67\%$; $p = 0.02$, one-sample t-test), whereas their estimated accuracy of the human advisor was slightly below its objective accuracy ($p_H^S = 49.66\% < p_H^O = 49.77\%$), with the difference being statistically insignificant ($p = 0.79$).

In the *Up-Down* treatment, subjects instead assigned a significantly lower prior subjective probability to the robo-advisor than to the human advisor (mean $p_R^S = 51.72\%$ versus mean $p_H^S = 53.91\%$; $p < 0.01$, paired t-test). Importantly, however, this difference reflects the substantially different objective accuracies of the two advisors. Relative to their respective objective probabilities, subjects again overestimated the robo-advisor's accuracy ($p_R^S = 51.72\% > p_R^O = 50.20\%$; $p < 0.01$, one-sample t-test), while slightly underestimating the human advisor's accuracy ($p_H^S = 53.91\% < p_H^O = 54.20\%$; $p = 0.57$).

To classify subjects' belief bias, we compare their estimated prior subjective probabilities with the corresponding objective probabilities measured in the pre-experiment. For each advisor $i \in \{H, R\}$, we classify subjects into three categories: positive belief bias ($p_i^S > p_i^O$), negative belief bias ($p_i^S < p_i^O$), and no belief bias ($p_i^S = p_i^O$).

As shown in Panel B of Table 5, a similar pattern emerges across the two treatments. In the *Odd-Even* treatment, 49.47% of subjects exhibit a positive belief bias toward the robo-advisor, compared with 42.63% toward the human advisor. The difference of 6.84 percentage points is not statistically significant ($p = 0.18$). Conversely, negative belief bias is less prevalent for the robo-advisor (27.37%) than for the human advisor (35.26%), a difference of 7.89 percentage points that is marginally significant ($p = 0.10$). The proportions exhibiting no belief bias are nearly identical across the two advisors (23.16% for the robo-advisor and 22.11% for the human advisor; $p = 0.81$).

The pattern is stronger in the *Up-Down* treatment. The proportion of subjects exhibiting positive belief bias is significantly higher for the robo-advisor than for the hu-

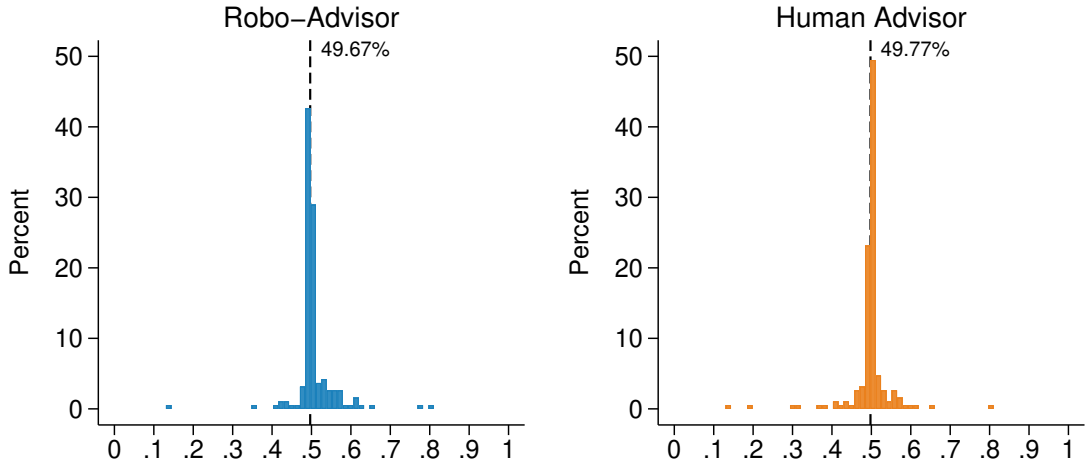

(a) *Odd-Even Treatment*

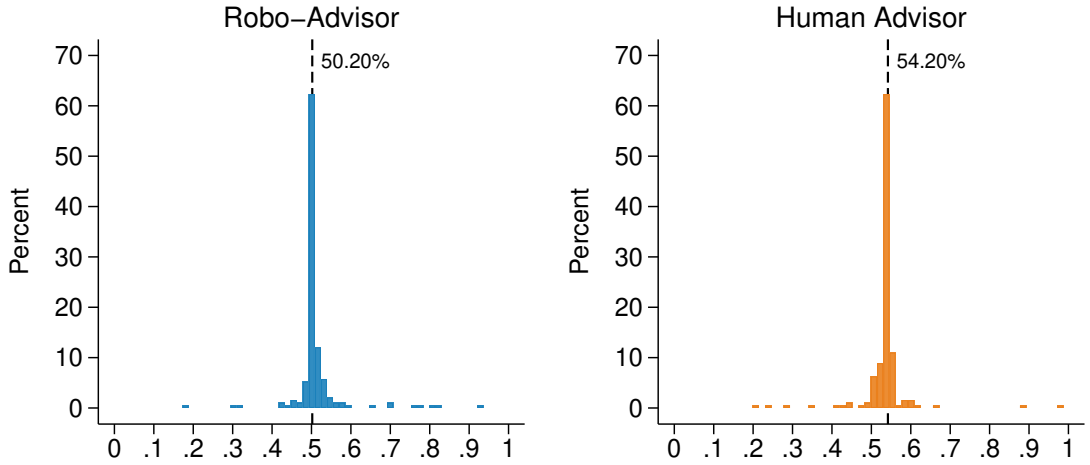

(b) *Up-Down Treatment*

Figure 7: Distribution of Estimated Prior Subjective Probability, p_i^S

man advisor (49.21% versus 36.13%), a difference of 13.09 percentage points ($p < 0.01$). In contrast, negative belief bias is significantly less prevalent for the robo-advisor than for the human advisor (33.51% versus 43.98%), a difference of 10.47 percentage points ($p = 0.04$). The proportions exhibiting no belief bias are similar across the two advisors (17.28% versus 19.90%; $p = 0.51$).

To further quantify the magnitude of belief bias, we compute the relative belief bias for each advisor as defined in Definition 2:

$$B_i = \frac{p_i^S - p_i^O}{p_i^O}, \qquad i \in \{H, R\}.$$

Table 5: Estimated Prior Subjective Probability and Belief Bias

	Mean / Proportion		Difference	p-value	
	Robot	Human		t-test	Wilcoxon
Panel A: Estimated Prior Subjective Probability					
Odd-Even	50.59%	49.66%	0.93%	0.00***	0.34
Up-Down	51.72%	53.91%	-2.18%	0.00***	0.00***
Panel B: Proportion of Subjects with Belief Bias					
Odd-Even					
Positive Belief Bias	49.47%	42.63%	6.84%	0.18	–
Negative Belief Bias	27.37%	35.26%	-7.89%	0.10*	–
No Belief Bias	23.16%	22.11%	1.05%	0.81	–
Up-Down					
Positive Belief Bias	49.21%	36.13%	13.09%	0.01***	–
Negative Belief Bias	33.51%	43.98%	-10.47%	0.04**	–
No Belief Bias	17.28%	19.90%	-2.62%	0.51	–

Notes: This table reports comparisons between the robo-advisor and the human advisor in the *Odd-Even* and *Up-Down* treatments. Panel A reports differences in estimated prior subjective probabilities using paired t-tests and Wilcoxon signed-rank tests. Panel B reports differences in the proportions of subjects exhibiting positive, negative, or no belief bias. For each advisor $i \in \{H, R\}$, positive belief bias is defined as $p_i^S > p_i^O$, negative belief bias as $p_i^S < p_i^O$, and no belief bias as $p_i^S = p_i^O$. Differences in proportions are tested using two-sample tests of proportions. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

In the *Odd-Even* treatment, subjects exhibit a positive mean belief bias toward the robo-advisor ($B_R = 1.86\%$) and a slightly negative mean belief bias toward the human advisor ($B_H = -0.21\%$). The mean difference, $B_R - B_H = 2.07\%$, is statistically significant ($p < 0.01$, paired t-test). Similarly, in the *Up-Down* treatment, the mean belief bias is positive for the robo-advisor ($B_R = 3.03\%$) and negative for the human advisor ($B_H = -0.54\%$). The corresponding difference, $B_R - B_H = 3.57\%$, is also statistically significant ($p < 0.01$, paired t-test).

Taken together, the results reveal a consistent asymmetry in subjects' beliefs across the two treatments. Although the robo-advisor does not necessarily receive a higher subjective probability in absolute terms, particularly in the *Up-Down* treatment where its objective accuracy is lower, subjects tend to evaluate the robo-advisor more favorably relative to its objective accuracy. Specifically, they overestimate the robo-advisor's accuracy while slightly underestimating the human advisor's accuracy. This pattern is consistent with a favorable belief bias toward robo-advice.

Result 5. In both the *Odd-Even* and *Up-Down* treatments, subjects exhibit a favorable belief bias toward the robo-advisor. Relative to objective accuracy, subjects tend to

overestimate the robo-advisor while slightly underestimating the human advisor. The resulting relative belief bias, $B_R - B_H$, is positive and significant in both treatments.

Source Preference. We next examine source preference using the estimated source-function parameters. Recall that β_i, $i \in \{R, H\}$, characterizes source preference under uncertainty, whereas β_i' characterizes source preference under risk after objective accuracy information is disclosed. Since a lower pessimism parameter implies a higher decision weight, a subject exhibits a source preference for the robo-advisor when $\beta_R < \beta_H$ under uncertainty and when $\beta_R' < \beta_H'$ under risk; the reverse inequalities characterize a preference for the human advisor.

Table 6: Source Preference for Robo-Advisor vs. Human Advisor

	Proportion		Difference	p-value
	First	Second		
Panel A: Odd-Even				
Across Sources				
Uncertainty	21.05% (Robot)	16.32% (Human)	4.74%	0.24
Risk	15.26% (Robot)	16.32% (Human)	−1.05%	0.78
Across Environments				
Robo-advisor	21.05% (Uncertainty)	15.26% (Risk)	5.79%	0.00***
Human advisor	16.32% (Uncertainty)	16.32% (Risk)	0.00%	1.00
Panel B: Up-Down				
Across Sources				
Uncertainty	58.12% (Robot)	13.61% (Human)	44.51%	0.00***
Risk	47.64% (Robot)	13.09% (Human)	34.55%	0.00***
Across Environments				
Robo-advisor	58.12% (Uncertainty)	47.64% (Risk)	10.47%	0.00***
Human advisor	13.61% (Uncertainty)	13.09% (Risk)	0.52%	0.66

Notes: This table reports the proportions of subjects classified as having a source preference for the robo-advisor or the human advisor. Under risk, source preference is determined by the estimated parameters $\widehat{\beta}_R'$ and $\widehat{\beta}_H'$. Under uncertainty, it is determined by $\widehat{\beta}_R = \widehat{\beta}_R' + \widehat{\delta}_R$ and $\widehat{\beta}_H = \widehat{\beta}_H' + \widehat{\delta}_H$. The corresponding source functions are evaluated at a common benchmark probability within each treatment, 0.4972 for *Odd-Even* and 0.5220 for *Up-Down*. A tolerance of 0.001 in decision weights is used to classify strict source preferences. "Across Sources" compares the proportions preferring the robo-advisor and the human advisor within a given information environment. "Across Environments" compares, for a given advisor, the proportion exhibiting a source preference under uncertainty with that under risk. All comparisons are based on paired t-tests at the individual level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 6 reports two complementary comparisons. First, we compare source preferences across advisors within each information environment. Second, we compare the prevalence of a given source preference between uncertainty and risk. The latter comparison captures

whether the classification of source preference changes when uncertainty about advisor accuracy is resolved.

Across Sources. In the *Odd-Even* treatment, we find no systematic preference for either source. Under uncertainty (i.e., before disclosure of accuracy information), 21.05% of subjects prefer the robo-advisor and 16.32% prefer the human advisor, a difference that is not statistically significant ($p = 0.24$). Under risk (i.e., after disclosure of accuracy information), the corresponding proportions are 15.26% and 16.32%, respectively, and are again statistically indistinguishable ($p = 0.78$). Thus, neither $\widehat{\beta}_R < \widehat{\beta}_H$ nor $\widehat{\beta}'_R < \widehat{\beta}'_H$ is systematically more prevalent than the reverse ordering.

The pattern is different in the *Up-Down* treatment. Under uncertainty, 58.12% of subjects prefer the robo-advisor, compared with 13.61% who prefer the human advisor, a difference of 44.51 percentage points ($p < 0.001$). Under risk, 47.64% prefer the robo-advisor and 13.09% prefer the human advisor, yielding a difference of 34.55 percentage points ($p < 0.001$). Hence, in *Up-Down*, the ordering $\widehat{\beta}_R < \widehat{\beta}_H$ is substantially more prevalent under uncertainty, and $\widehat{\beta}'_R < \widehat{\beta}'_H$ remains substantially more prevalent after accuracy information is disclosed.

Across Information Environments. Information disclosure affects the prevalence of robo-advisor preference but not human-advisor preference. In *Odd-Even*, the proportion preferring the robo-advisor decreases from 21.05% under uncertainty to 15.26% under risk, a decline of 5.79 percentage points ($p < 0.001$). In contrast, the proportion preferring the human advisor is unchanged at 16.32%.

The same pattern is stronger in *Up-Down*. The proportion preferring the robo-advisor decreases from 58.12% under uncertainty to 47.64% under risk, a decline of 10.47 percentage points ($p < 0.001$), whereas the proportion preferring the human advisor changes only from 13.61% to 13.09%, a difference that is not statistically significant ($p = 0.66$). Thus, resolving uncertainty about advisor accuracy reduces the prevalence of robo-advisor preference in both treatments, while leaving the prevalence of human-advisor preference essentially unchanged.

Result 6. Source preference differs across tasks and information environments. In *Odd-Even*, there is no systematic preference for either advisor, although robo-advisor preference is more prevalent under uncertainty than under risk. In *Up-Down*, subjects systematically prefer the robo-advisor in both environments, with a stronger preference under uncertainty. Human-advisor preference remains stable across environments in both treatments.

5 Conclusion

We experimentally investigate the roles of belief bias and source preference in shaping individuals' preferences for investment advice from a robo-advisor versus a human advisor. Our experiment extends the source preference framework by incorporating a dynamic belief–taste decomposition. Departing from prior literature (Fox and Tversky, 1995; Chew et al., 2008; Abdellaoui et al., 2011; Li, 2011; Chew et al., 2017; Li et al., 2018), which relies on static measures, our design provides feedback on the actual performance of both advisors. This allows us to disentangle belief bias (i.e., incorrect priors about advisor competence) from source preference, defined as a preference for the advisor independent of performance.

Our findings show that both belief bias and source preference contribute to preferences for robo-advisors. Before accuracy information was disclosed, subjects were significantly more likely to prefer the robo-advisor over the human advisor, even though the investments were objectively identical. After subjects learned that the two advisors performed similarly in the *Odd-Even* task, the proportion preferring the robo-advisor declined, while the proportion preferring the human advisor increased. This shift indicates that, for some subjects, the initial preference for the robo-advisor reflected belief bias that was corrected by performance feedback.

At the same time, a substantial proportion of subjects continued to pay a premium for robo-advice even after accuracy disclosure. This persistent preference cannot be explained by beliefs about performance and therefore identifies source preference: a preference for the robo-advisor itself, independent of its competence. Thus, while belief bias helps explain the initial overvaluation of robo-advisors, a larger component of subjects' preference for robo-advice appears to reflect source preference. We observe qualitatively similar patterns in the *Up-Down* treatment, demonstrating the robustness of these findings across contexts.

Another finding of our study is that a significant proportion of subjects have a stronger preference for robo-advisor under uncertainty, suggesting that they believe that robots are more competent than human advisor, supporting the competence hypothesis of (Heath and Tversky, 1991). Further, in the *Odd-Even* treatment, source preferences are stable, while that is not the case in the *Up-Down* treatment. This suggests that when the advisor's quality is ambiguous (as in the *Up-Down* treatment), people's preference is unstable. Thus, when the disclosed accuracy is new information, and people update their beliefs about how good each source is, and their preferences shift. On the other hand, in the *Odd-Even* treatment, disclosed accuracy tells people little they did not already expect. Thus, the result implies that the preference for robo-advisor will be less stable when there is uncertainty.

Our findings have two main policy implications. First, mandatory disclosure of robo-advisor performance can improve welfare by correcting belief bias, as the shift in valuations after disclosure shows that part of the initial robo-premium reflects mistaken beliefs about algorithmic competence. Second, the presence of source preference suggests that demand for robo-advice may persist even when performance information is fully available. In our experiment, the source has no instrumental value. The bet's payoff depends only on the advisor's accuracy, and accuracy is known and slightly lower for the robo-advisor. Whether performance-independent source preference is welfare-reducing depends on whether the source carries value beyond performance. In our design, it does not, so the residual premium represents foregone expected payoff, though in field settings such preferences may reflect genuine consumption value. Disclosure is therefore necessary but not sufficient to align reliance on robo-advice with advisor performance, and complementary interventions, such as changes in design, communication, or accountability, may be needed.

Our findings suggest several directions for future research. First, while our design provides a single episode of performance feedback, many real-world investment decisions involve repeated interactions. Future work could examine whether repeated exposure to equal performance eventually eliminates source preference, or whether such preferences persist and perhaps even widen over time. Second, although Stage 2 of our experiment provides full information about advisors' accuracy, real-world environments often involve

noisy or incomplete feedback. Future research could study how source preference operates when performance signals leave residual uncertainty about advisors' true competence, and whether ambiguity attenuates the correction of belief bias or increases reliance on source preference.

References

Abdellaoui, M., A. Baillon, L. Placido, and P. P. Wakker (2011): "The rich domain of uncertainty: Source functions and their experimental implementation," *American Economic Review*, 101, 695–723.

Almog, D., R. Gauriot, L. Page, and D. Martin (2024): "Human Responses to Algorithmic Correction: Evidence from Centre Court," *arXiv preprint arXiv:2401.16754*.

Almog, D., L. Lippman, and D. Martin (2026): "When an AI Judges Your Work: The Hidden Costs of Algorithmic Assessment," *arXiv preprint arXiv:2603.02076*.

Ambuehl, S., B. D. Bernheim, F. Ersoy, and D. Harris (2025): "Peer Advice on Financial Decisions: A case of the blind leading the blind?" *Review of Economics and Statistics*, 107, 240–255.

Autor, D., A. Caplin, D. J. Martin, and P. Marx (2025): "Misaligned by Design: Incentive Failures in Machine Learning," *Working Paper*.

Back, C., S. Morana, and M. Spann (2023): "When do robo-advisors make us better investors? The impact of social design elements on investor behavior," *Journal of Behavioral and Experimental Economics*, 103, 101984.

Baillon, A., H. Bleichrodt, C. Li, and P. P. Wakker (2025): "Source theory: A tractable and positive ambiguity theory," *Management Science*, 71, 8767–8782.

Brenner, L. and T. Meyll (2020): "Robo-advisors: a substitute for human financial advice?" *Journal of Behavioral and Experimental Finance*, 25, 100275.

Campbell, J. L., H. Ham, Z. Lu, and K. Wood (2026): "Expectations matter: When (not) to use machine learning earnings forecasts," *Management Science*, 72, 8863–8887.

Caplin, A., D. Deming, S. Li, D. Martin, P. Marx, B. Weidmann, and K. J. Ye (2026): "The ABCs of who benefits from working with AI: Ability, beliefs, and calibration," *Management Science*, 72, 5843–5852.

Chen, D. L., M. Schonger, and C. Wickens (2016): "oTree—An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Chen, Y. and S. Zhong (2024): "Source Dependence in Effort Provision," *International Economic Review*, 65, 1499–1517.

Chew, S. H. and K. K. Li (2021): "The moral investor: Sin stock aversion and virtue stock affinity," *Available at SSRN 3773971*.

Chew, S. H., K. K. Li, R. Chark, and S. Zhong (2008): "Source preference and ambiguity aversion: Models and evidence from behavioral and neuroimaging experiments," in *Neuroeconomics*, Emerald Group Publishing Limited, 179–201.

Chew, S. H., B. Miao, and S. Zhong (2017): "Partial ambiguity," *Econometrica*, 85, 1239–1260.

Chew, S. H. and J. S. Sagi (2008): "Small worlds: Modeling attitudes toward sources of uncertainty," *Journal of Economic Theory*, 139, 1–24.

Choi, H., J. Kim, M. Kovach, K.-M. Lee, E. Shin, and H. Tzavellas (2026): "Do People Follow AI Advice? Evidence from a Pension Portfolio Choice Experiment," *arXiv preprint arXiv:2608.11371*.

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

Fernandes, D., J. G. Lynch Jr, and R. G. Netemeyer (2014): "Financial literacy, financial education, and downstream financial behaviors," *Management Science*, 60, 1861–1883.

Fox, C. R. and A. Tversky (1995): "Ambiguity aversion and comparative ignorance," *The Quarterly Journal of Economics*, 110, 585–603.

He, H. and K. K. Li (2026): "Illusion on Causality: An Experimental Study," *Working Paper*.

Heath, C. and A. Tversky (1991): "Preference and belief: Ambiguity and competence in choice under uncertainty," *Journal of Risk and Uncertainty*, 4, 5–28.

Hodge, F. D., K. I. Mendoza, and R. K. Sinha (2021): "The effect of humanizing robo-advisors on investor judgments," *Contemporary Accounting Research*, 38, 770–792.

Holzmeister, F., M. Holmén, M. Kirchler, M. Stefan, and E. Wengström (2023): "Delegation decisions in finance," *Management Science*, 69, 4828–4844.

Kovach, M., D. Martin, and G. Tserenjigmid (2026): "Learning from an Unknown DGP: Experimental Evidence on Belief Updating with AI Recommendations," *arXiv preprint arXiv:2607.10460*.

Li, K. K. (2011): "Preference towards control in risk taking: Control, no control, or randomize?" *Journal of Risk and Uncertainty*, 43, 39–63.

Li, Z., J. Müller, P. P. Wakker, and T. V. Wang (2018): "The rich domain of ambiguity explored," *Management Science*, 64, 3227–3240.

Linnainmaa, J. T., B. T. Melzer, and A. Previtero (2021): "The misguided beliefs of financial advisors," *The Journal of Finance*, 76, 587–621.

Machina, M. J. (2004): "Almost-objective uncertainty," *Economic Theory*, 24, 1–54.

Magesh, V., F. Surani, M. Dahl, M. Suzgun, C. D. Manning, and D. E. Ho (2025): "Hallucination-free? Assessing the reliability of leading AI legal research tools," *Journal of Empirical Legal Studies*, 22, 216–242.

Pouget, S., J. Sauvagnat, and S. Villeneuve (2017): "A mind is a terrible thing to change: confirmatory bias in financial markets," *The Review of Financial Studies*, 30, 2066–2109.

Savage, L. J. (1954): *The Foundations of Statistics*, New York: Wiley.

Schwalbe, N. and B. Wahl (2020): "Artificial intelligence and the future of global health," *The Lancet*, 395, 1579–1586.

Shanmuganathan, M. (2020): "Behavioural finance in an era of artificial intelligence: Longitudinal case study of robo-advisors in investment decisions," *Journal of Behavioral and Experimental Finance*, 27, 100297.

Statman, M., S. Thorley, and K. Vorkink (2006): "Investor overconfidence and trading volume," *The Review of Financial Studies*, 19, 1531–1565.

Tse, T. T. K., N. Hanaki, and B. Mao (2024): "Beware the performance of an algorithm before relying on it: Evidence from a stock price forecasting experiment," *Journal of Economic Psychology*, 102, 102727.

A Online Appendix

A.1 Estimation

We estimate subjects' subjective probabilities and source-function parameters using the structural framework introduced in the main text. For each advisor $i \in \{H, R\}$, let

$$p_i^O = P(E_i)$$

denote the objective probability that advisor i makes a correct prediction, and let

$$p_i^S = P_i(E_i)$$

denote the subject's subjective probability of the same event before the advisor's accuracy is disclosed. Thus, p_i^O corresponds to the risk environment after information disclosure, whereas p_i^S corresponds to the uncertainty environment before information disclosure.

Following the main-text specification, we use the source function

$$W(p; \alpha, \beta) = \exp\left[-\beta(-\ln p)^\alpha\right].$$

The source functions in the two information environments are therefore

$$w_i(p) = W(p; \alpha, \beta_i)$$

under uncertainty and

$$w_i'(p) = W(p; \alpha, \beta_i')$$

under risk. Here, β_i is the pessimism parameter under uncertainty and β_i' is the corresponding parameter under risk. The common parameter α captures likelihood insensitivity. Consistent with the theoretical framework, we impose a common α across advisors and information environments, while allowing β_H, β_R, β_H', and β_R' to differ.

Each subject reports certainty equivalents for five payoff pairs,

$$(x, y) \in \{(10, 0), (15, 0), (25, 0), (20, 10), (25, 15)\},$$

for both advisors and both probability conditions. Let $j = 1, \ldots, 5$ index the payoff pairs.

We assume the normalized power utility function

$$u(x) = \left(\frac{x}{25}\right)^{\rho},$$

where ρ governs utility curvature. The model-implied certainty equivalent for advisor i, payoff pair j, and probability condition $k \in \{S, O\}$ is

$$\widehat{CE}_{ij}^{k} = 25\left[w_i^k(p_i^k)\left(\frac{x_j}{25}\right)^{\rho} + \left(1 - w_i^k(p_i^k)\right)\left(\frac{y_j}{25}\right)^{\rho}\right]^{1/\rho},$$

where

$$w_i^S(p) = W(p; \alpha, \beta_i), \qquad w_i^O(p) = W(p; \alpha, \beta_i'),$$

and

$$p_i^k = \begin{cases} p_i^S, & k = S, \\ p_i^O, & k = O. \end{cases}$$

We estimate the model separately for each subject using two stages.

Stage 1: Risk-state utility and source-function parameters. In the first stage, we use only the objective-probability valuation tasks. The probabilities p_R^O and p_H^O are known from our pre-experiment. We estimate the utility-curvature parameter ρ, the common likelihood-insensitivity parameter α, and the risk-state source-function parameters β_R' and β_H' by minimizing

$$\sum_j \left[CE_{Rj}^O - \widehat{CE}_{Rj}^O\left(p_R^O; \rho, \alpha, \beta_R'\right)\right]^2 + \sum_j \left[CE_{Hj}^O - \widehat{CE}_{Hj}^O\left(p_H^O; \rho, \alpha, \beta_H'\right)\right]^2.$$

The objective probabilities are treatment-specific. In the *Odd-Even* treatment,

$$p_R^O = 0.4967, \qquad p_H^O = 0.4977,$$

whereas in the *Up-Down* treatment,

$$p_R^O = 0.5020, \qquad p_H^O = 0.5420.$$

Because the number of observations at the subject level is small and the model is nonlinear, we use a grid search. The search grids are

$$\rho \in \{0.2, 0.4, \ldots, 3.0\},$$

$$\alpha \in \{0.2, 0.4, \ldots, 2.6\},$$

and

$$\beta_i' \in \{0.1, 0.5, 1.0, 1.5, \ldots, 5.0, 6.0, 7.0, 8.0, 9.0, 10.0, 12.0, 15.0\}.$$

For each candidate pair (ρ, α), we choose β_R' and β_H' separately to minimize the advisor-specific residual sums of squares. We then select the combination

$$(\widehat{\rho}, \widehat{\alpha}, \widehat{\beta}_R', \widehat{\beta}_H')$$

that minimizes the combined objective across the two advisors.

Stage 2: Subjective probabilities and uncertainty-state source functions. In the second stage, we hold

$$\widehat{\rho}, \quad \widehat{\alpha}, \quad \widehat{\beta}_R', \quad \widehat{\beta}_H'$$

fixed and estimate the subjective probabilities and the changes in source weighting from the subjective-probability valuation tasks.

The key feature of the estimation is that we allow the source-function parameter to differ between the uncertainty and risk environments. For each advisor $i \in \{H, R\}$, we specify

$$\beta_i = \beta_i' + \delta_i,$$

where β_i' is the risk-state pessimism parameter estimated in Stage 1, β_i is the uncertainty-state pessimism parameter, and δ_i captures the advisor-specific change in source weighting between the two environments. Thus, the source function used for the subjective-probability tasks is

$$w_i(p_i^S) = W(p_i^S; \alpha, \beta_i) = W(p_i^S; \alpha, \beta_i' + \delta_i).$$

This parameterization directly corresponds to the theoretical framework in the main text. In the notation of the main text, β_i' characterizes source preference under risk, after advisor accuracy has been disclosed, whereas β_i characterizes source preference under uncertainty, before accuracy information has been disclosed. The stability restriction considered in the theoretical discussion is the nested restriction

$$\delta_i = 0,$$

under which $\beta_i = \beta_i'$ and the source function is unchanged across the two information environments.

To keep estimated subjective probabilities within the unit interval and to avoid numerical solutions at the boundary, we set

$$p_L = 0.001, \qquad p_U = 0.999.$$

For each advisor, define

$$a_i = \log\left(\frac{p_i^O - p_L}{p_U - p_i^O}\right).$$

We parameterize the subjective probability as

$$p_i^S(\theta_i) = p_L + (p_U - p_L)\,\text{logit}^{-1}(a_i + \theta_i).$$

Consequently,

$$\theta_i = 0 \quad \Longrightarrow \quad p_i^S = p_i^O,$$

while positive and negative values of θ_i imply subjective probabilities above and below the objective probability, respectively.

For each advisor, we jointly estimate θ_i and δ_i by minimizing

$$\mathcal{L}_i(\theta_i, \delta_i) = \sum_j \left[CE_{ij}^S - \widehat{CE}_{ij}^S\left(p_i^S(\theta_i); \widehat{\rho}, \widehat{\alpha}, \widehat{\beta}_i' + \delta_i\right)\right]^2 + \lambda_i\theta_i^2 + \lambda_{\delta,i}\delta_i^2.$$

The first penalty regularizes deviations of the subjective probability from the objective probability. The second penalty regularizes changes in the source function toward the stability benchmark $\delta_i = 0$.

Both penalty weights are determined by the distance between matched subjective- and objective-probability certainty equivalents:

$$\lambda_i = 25\exp\left(-\overline{|CE_i^S - CE_i^O|}\right),$$

and

$$\lambda_{\delta,i} = 25\exp\left(-\overline{|CE_i^S - CE_i^O|}\right),$$

where the bar denotes the average across matched payoff pairs for advisor i. Thus, when the subjective- and objective-probability certainty equivalents are nearly identical,

both the subjective probability and the source-function shift are more strongly anchored toward their stability benchmarks. When the two sets of certainty equivalents differ substantially, the data are allowed to generate larger deviations in both dimensions.

The two advisors are estimated separately. Conditional on the first-stage estimates, the objective for advisor i depends only on (θ_i, δ_i). We search over

$$\delta_i \in [-5, 5]$$

in increments of 0.1. Conditional on each candidate value of δ_i, we search over

$$\theta_i \in [-12, 12]$$

using a coarse grid with step size 0.05. We then refine the search within a neighborhood of ± 0.10 around the best coarse-grid value using step size 0.005. The selected pair

$$(\widehat{\theta}_i, \widehat{\delta}_i)$$

minimizes the advisor-specific penalized objective.

In the implementation, the effective uncertainty-state pessimism parameter is constrained to remain positive:

$$\widetilde{\beta}_i = \max\left\{0.001, \widehat{\beta}'_i + \delta_i\right\}.$$

The source function used in the subjective-probability valuation tasks is therefore evaluated at

$$W\left(p_i^S(\theta_i); \widehat{\alpha}, \widetilde{\beta}_i\right).$$

The estimated uncertainty-state parameter is reported as

$$\widehat{\beta}_i = \widetilde{\beta}_i,$$

while $\widehat{\delta}_i$ records the estimated change relative to the risk-state parameter $\widehat{\beta}'_i$.

This procedure estimates the parameters required for both types of source preference defined in the main text. Source preference under risk is characterized by the comparison between

$$\widehat{\beta}'_R \quad \text{and} \quad \widehat{\beta}'_H,$$

which are estimated from the objective-probability valuation tasks. Source preference under uncertainty is characterized by the comparison between

$$\widehat{\beta}_R = \widehat{\beta}'_R + \widehat{\delta}_R \quad \text{and} \quad \widehat{\beta}_H = \widehat{\beta}'_H + \widehat{\delta}_H,$$

which are obtained from the subjective-probability valuation tasks.

Accordingly, the difference between source preference under uncertainty and source preference under risk may arise from two components: the baseline difference in the risk-state parameters β'_R and β'_H, and the advisor-specific changes in source weighting δ_R and δ_H. Because the shifts are estimated separately, the procedure does not impose that the source function change by the same amount for the human and robo-advisor.

The source-function stability hypothesis corresponds to the nested restriction

$$\delta_R = \delta_H = 0.$$

Therefore, the estimated values of $\widehat{\delta}_R$ and $\widehat{\delta}_H$ provide a direct diagnostic of whether the source functions are stable across the uncertainty and risk environments. Values close to zero are consistent with the stability assumption, whereas systematic deviations from zero indicate that the source-function parameters differ between the two information environments.

A.2 Additional Tables

Table A1: Valuations for Robo vs. Human Advisor after Information Update

	Mean Valuation		Difference	p-value	
	Robot	Human		t-test	Wilcoxon
Odd-Even	10.14	10.31	-0.17	0.21	0.65
Up-Down	9.94	10.38	-0.44	0.01***	0.01**

Notes: The p-values are based on paired t-tests and Wilcoxon signed-rank tests. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A2: Proportion Preferring Robo vs. Human Advisor after Information Update

	Proportion		Difference	p-value
	Robot	Human		
Odd-Even	26.84%	28.42%	-1.58%	0.73
Up-Down	26.18%	39.79%	-13.61%	0.00***

Notes: The p-values are based on two-sample t-tests of proportions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A3: Preference Transition in *Odd-Even* Treatment

	Revised Preference			
Initial Preference	Robot	Human	Indifference	Total
Robot	47 (24.74%)	11 (5.79%)	11 (5.79%)	69
Human	2 (1.05%)	39 (20.53%)	3 (1.58%)	44
Indifference	2 (1.05%)	4 (2.11%)	71 (37.37%)	77
Total	51	54	85	190

Table A4: Preference Transition in *Up-Down* Treatment

	Revised Preference			
Initial Preference	Robot	Human	Indifference	Total
Robot	44 (23.04%)	18 (9.42%)	6 (3.14%)	68
Human	3 (1.57%)	39 (20.42%)	3 (1.57%)	45
Indifference	3 (1.57%)	19 (9.95%)	56 (29.32%)	78
Total	50	76	65	191

A.3 Regression Analyses

The results reported in Section 4 are also supported by the OLS regressions reported in Table A5 and the Probit regressions reported in Table A6.

Table A5: Determinants of Valuations

	Odd-Even		Up-Down	
	Robot	Human	Robot	Human
	(1)	(2)	(3)	(4)
Information Update	-0.333(***)	0.020	-0.473(***)	0.058
	(0.103)	(0.111)	(0.142)	(0.126)
Controls	YES	YES	YES	YES
Constant	5.424(***)	5.389(***)	6.391(***)	7.098(***)
	(1.676)	(1.686)	(1.843)	(1.830)
Observations	1900	1900	1910	1910
R^2	0.579	0.578	0.571	0.551

Notes: OLS regressions using data from the *Odd-Even* and *Up-Down* treatments. The dependent variable is Valuation, which is subjects' elicited valuation for receiving advice from the advisor. Information Update is a dummy variable equal to 1 if the valuation is elicited after subjects are informed of the advisors' prediction accuracy, and 0 otherwise. Columns (1) and (2) report estimates from the *Odd-Even* treatment, and Columns (3) and (4) from the *Up-Down* treatment. Within each treatment, the first column reports regressions using valuations for the robo-advisor, and the second column reports regressions using valuations for the human advisor. All regressions include controls for valuation-type fixed effects as well as subjects' demographic characteristics (gender, age, income, and education). In addition, we control for subjects' self-reported familiarity with robo-advisors and human advisors, as well as their prior experience with each type of advisor (dummy variables equal to 1 if the subject reports any prior experience). Standard errors in parentheses are clustered at the subject level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

As reported in Table A5, the coefficient on the information-update dummy is negative and highly significant for the robo-advisor in both treatments (*Odd-Even*: -0.333; *Up-Down*: -0.473), confirming that disclosure of accuracy information significantly reduces valuations for robo-advice. For example, in the *Odd-Even* treatment, subjects' valuations for the robo-advisor decrease by 0.333 ECU on average after the information update. In contrast, the coefficient is statistically insignificant for the human advisor in both treatments. These regression results support the changes in valuations documented in Section 4.

Table A6 reports the corresponding Probit regressions, with marginal effects reported for the probability of preferring each advisor. The dependent variable is a dummy equal to

Table A6: Determinants of Proportion Preferring Advisors

	Odd-Even		Up-Down	
	Robot	Human	Robot	Human
	(1)	(2)	(3)	(4)
Information Update	-0.041(***)	0.038(***)	-0.043(***)	0.115(***)
	(0.011)	(0.013)	(0.014)	(0.019)
Controls	YES	YES	YES	YES
Observations	1900	1900	1910	1910

Notes: Probit regressions using data from the *Odd-Even* and *Up-Down* treatments. Reported coefficients are marginal effects. The dependent variable is Prefer Robo-Advisor in Columns (1) and (3), and Prefer Human Advisor in Columns (2) and (4). The dummy variable Prefer Robo-Advisor equals 1 if a subject's average valuation for the robo-advisor is strictly higher than that for the human advisor, and 0 otherwise. The dummy variable Prefer Human Advisor equals 1 if a subject's valuation for the human advisor is strictly higher than that for the robo-advisor for a given lottery, and 0 otherwise. Columns (1) and (2) report estimates from the *Odd-Even* treatment, and Columns (3) and (4) from the *Up-Down* treatment. All regressions include lottery-type fixed effects as well as subjects' demographic characteristics (gender, age, income, and education). In addition, we control for subjects' self-reported familiarity with robo-advisors and human advisors, and their prior experience with each type of advisor (dummy variables equal to 1 if the subject reports any prior experience). Standard errors in parentheses are clustered at the subject level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

1 if a subject's valuation for the robo-advisor (human advisor) is strictly higher than that for the human advisor (robo-advisor) for a given lottery. The structure mirrors Table A5. The coefficient on the information-update dummy is negative and highly significant for the probability of preferring the robo-advisor in both treatments (*Odd-Even*: -0.041; *Up-Down*: -0.043), while it is positive and highly significant for the probability of preferring the human advisor (*Odd-Even*: 0.038; *Up-Down*: 0.115). For example, in the *Odd-Even* treatment, the probability of preferring the robo-advisor decreases by approximately 4.1 percentage points after the information update, while the probability of preferring the human advisor increases by approximately 3.8 percentage points. These results are consistent with the observed changes in preferences and confirm that the information update shifts subjects' choices away from the robo-advisor and toward the human advisor or indifference.[9]

[9]Additionally, in the regressions reported in Tables A5 and A6, we find no significant relationships between subjects' familiarity, prior experience, or knowledge of financial technologies and their valuations of the advisors.

B Experimental Instructions: Pre-Experiment

B.1 *Odd-Even* Treatment

(The First Page)[10]

Welcome to our Decision-Making Study!

The entire study, including tasks and a questionnaire, will take approximately **8** minutes to complete.

In addition to a participation fee, you have the opportunity to earn a bonus based on your decisions. After completing the experiment, one task involving payment will be randomly drawn, and your bonus will be determined by the decision you made in that task.

Your final payoff will include both your participation fee and the bonus from the randomly drawn task.

Please take your time to carefully read and thoughtfully respond to all the questions. All decisions and personal information in this study will remain completely anonymous and strictly confidential. Thank you for your participation!

Please wait 10 seconds...[11]

(Continue)

[10]Each subject was presented with a consent form at the beginning of the experiment, and only those who accepted the consent form could proceed with the study. The setting is the same for other two treatments (i.e., *Up-Down* and *Numbers*).

[11]We set a timer here to ensure that subjects had some time to read the instructions carefully; similar timers were used throughout all our experiments for the same purpose.

(Page 2)

Instructions

In this task, you will be provided with historical data on a stock's closing prices from Day 1 to Day 30 in an undisclosed year, without revealing the stock's ticker symbol. You will make predictions on whether the last digit of the stock's closing price was odd or even for each trading day from Day 31 to Day 60 in the same year.

The Stock

The stock is a real, publicly traded stock randomly selected from the stock exchange. This stock will remain the same throughout today's experiment. The 60 consecutive trading days are labeled as Day 1 to Day 60. The first trading day is Day 1, the second is Day 2, and so on.

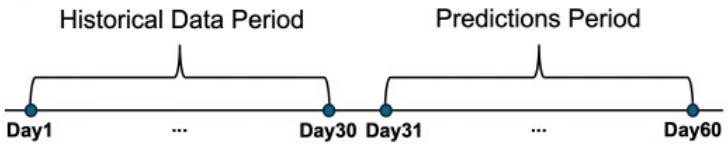

The Decision Task

You will make predictions on **whether the last digit of the stock's closing price was odd or even** for each trading day from Day 31 to Day 60. You will choose one option (either odd or even) and if your choice matches the actual digit's parity (i.e., odd or even), you will win $2; otherwise, you will win nothing.

An example of what you will see in the main tasks: in the following, you will make the prediction on whether the last digit of the stock's closing price was odd or even on Day 31.

Please make the prediction on whether the last digit of the stock's closing price was odd or even on **Day 31**. Please choose one option.

Odd Even

Suppose you choose Odd, and the actual digit was Odd, you will win $2.

Suppose you choose Odd, and the actual digit was Even, you will win nothing.

Your Payoff

After completing the experiment, one of your predictions from Day 31 to Day 60 will be randomly drawn for the bonus payment. If your prediction for the drawn day is correct, you will earn **$2**; otherwise, you will receive no bonus.

Historical Data

The stock's closing prices from Day 1 to Day 30 are shown as follows:

Historical Data	
Date	**Closing Prices**
Day 1	75.68
Day 2	73.23
Day 3	72.26
Day 4	66.54
Day 5	71.33
Day 6	68.86
Day 7	62.06
Day 8	69.49
Day 9	60.55
Day 10	63.22
Day 11	61.67
Day 12	61.19
Day 13	57.31
Day 14	56.09
Day 15	61.72
Day 16	61.38
Day 17	64.61
Day 18	61.94
Day 19	63.70
Day 20	63.57
Day 21	60.23
Day 22	61.23
Day 23	60.35
Day 24	65.62
Day 25	64.86
Day 26	66.52
Day 27	67.00
Day 28	68.31
Day 29	71.76
Day 30	71.11

Historical data is available on each page where you make your predictions. Before the main tasks, we would like you to complete two practice tasks to ensure that you fully understand the rules. Please read the instructions above carefully and click the "Continue" button to begin the practice section.

You will be able to continue after **30** seconds. Please read the instructions carefully.

(Continue)

(Page 3)

Practice Task 1

Historical Data

The stock's closing prices from Day 1 to Day 30 are shown as follows:

Historical Data	
Date	**Closing Prices**
Day 1	75.68
Day 2	73.23
Day 3	72.26
Day 4	66.54
Day 5	71.33
Day 6	68.86
Day 7	62.06
Day 8	69.49
Day 9	60.55
Day 10	63.22
Day 11	61.67
Day 12	61.19
Day 13	57.31
Day 14	56.09
Day 15	61.72
Day 16	61.38
Day 17	64.61
Day 18	61.94
Day 19	63.70
Day 20	63.57
Day 21	60.23
Day 22	61.23
Day 23	60.35
Day 24	65.62
Day 25	64.86
Day 26	66.52
Day 27	67.00
Day 28	68.31
Day 29	71.76
Day 30	71.11

Please make the prediction on whether the last digit of the stock's closing price was odd or even on **Day XX**. Please choose one option.

Odd Even

Note: you can click the following button to read the instructions for this task.

[Instructions]

(Page 4)

Practice Task 2

Historical Data

The stock's closing prices from Day 1 to Day 30 are shown as follows:

Historical Data	
Date	**Closing Prices**
Day 1	75.68
Day 2	73.23
Day 3	72.26
Day 4	66.54
Day 5	71.33
Day 6	68.86
Day 7	62.06
Day 8	69.49
Day 9	60.55
Day 10	63.22
Day 11	61.67
Day 12	61.19
Day 13	57.31
Day 14	56.09
Day 15	61.72
Day 16	61.38
Day 17	64.61
Day 18	61.94
Day 19	63.70
Day 20	63.57
Day 21	60.23
Day 22	61.23
Day 23	60.35
Day 24	65.62
Day 25	64.86
Day 26	66.52
Day 27	67.00
Day 28	68.31
Day 29	71.76
Day 30	71.11

Please make the prediction on whether the last digit of the stock's closing price was odd or even on **Day XX**. Please choose one option.

Odd Even

Note: you can click the following button to read the instructions for this task.

[Instructions]

(Page 5)

Practice Tasks Completed

You have completed the practice tasks. Please click the "Continue" button to begin the main tasks.

(Continue)

(Pages 6-35)[12]

Main Task - Day 31

Historical Data

The stock's closing prices from Day 1 to Day 30 are shown as follows:

Historical Data	
Date	**Closing Prices**
Day 1	75.68
Day 2	73.23
Day 3	72.26
Day 4	66.54
Day 5	71.33
Day 6	68.86
Day 7	62.06
Day 8	69.49
Day 9	60.55
Day 10	63.22
Day 11	61.67
Day 12	61.19
Day 13	57.31
Day 14	56.09
Day 15	61.72
Day 16	61.38
Day 17	64.61
Day 18	61.94
Day 19	63.70
Day 20	63.57
Day 21	60.23
Day 22	61.23
Day 23	60.35
Day 24	65.62
Day 25	64.86
Day 26	66.52
Day 27	67.00
Day 28	68.31
Day 29	71.76
Day 30	71.11

Please make the prediction on whether the last digit of the stock's closing price was odd or even on **Day 31**. Please choose one option.

Odd Even

Note: you can click the following button to read the instructions for this task.

[Instructions]

[12]In the current page, we use the main task on Day 31 as an example. The subsequent 29 pages only involves changes from Day 31 to Day 32, and so on, up to Day 60. The setting is the same for other two treatments (i.e., *Up-Down* and *Numbers*).

(Page 36)

Questionnaire

This section includes a questionnaire that we would like you to complete. Please answer the questions thoughtfully and accurately. Your responses will be kept confidential and anonymous. Thank you!

(Continue)

(Page 37)

Questionnaire

What is your gender?

○ Male

○ Female

(Continue)

(Page 38)

Questionnaire

What is your age?

_____

(Continue)

(Page 39)

Questionnaire

What is your annual income?

○ Less than $10,000

○ Between $10,000 - $19,999

○ Between $20,000 - $29,999

○ Between $30,000 - $39,999

○ Between $40,000 - $49,999

○ Between $50,000 - $59,999

○ Between $60,000 - $69,999

○ Between $70,000 - $79,999

○ Between $80,000 - $89,999

○ Between $90,000 - $99,999

○ Between $100,000 - $149,999

○ More than $150,000

(Continue)

(Page 40)

Questionnaire

What is the highest degree or level of school you have completed?

○ Some elementary school or below

○ Elementary school

○ Middle school

○ High school or equivalent

○ Some college/associate degree

○ Bachelor's degree

○ Master's degree

○ Professional or PhD

(Continue)

(Page 41)

Questionnaire

Are you currently...?

○ Employed for wages

○ Self-employed

○ Out of work and looking for work

○ Out of work but not currently looking for work

○ A homemaker

○ A student

○ Military

○ Retired

○ Unable to work

(Continue)

(The Last Page)

Experiment Finished

Thank you for participating in this experiment!

Please do not forget to click the "Continue and confirm your participation" button in the lower right corner of this page. Only in this way will your submission be successfully received, and only the successful submission will be approved and paid.

(Continue and confirm your participation)

B.2 *Up-Down* Treatment

(The First Page)

Welcome to our Decision-Making Study!

The entire study, including tasks and a questionnaire, will take approximately **8** minutes to complete.

In addition to a participation fee, you have the opportunity to earn a bonus based on your decisions. After completing the experiment, one task involving payment will be randomly drawn, and your bonus will be determined by the decision you made in that task.

Your final payoff will include both your participation fee and the bonus from the randomly drawn task.

Please take your time to carefully read and thoughtfully respond to all the questions. All decisions and personal information in this study will remain completely anonymous and strictly confidential. Thank you for your participation!

Please wait 10 seconds...

(Continue)

(Page 2)

Instructions

In this task, you will be provided with historical data on a stock's closing price changes from Day 1 to Day 30 compared to the previous trading day in an undisclosed year, without revealing the stock's ticker symbol. You will make predictions on whether the stock's closing price was up or down compared to the previous trading day for each trading day from Day 31 to Day 60 in the same year.

The Stock

The stock is a real, publicly traded stock randomly selected from the stock exchange. This stock will remain the same throughout today's experiment. The 60 consecutive trading days are labeled as Day 1 to Day 60. The first trading day is Day 1, the second is Day 2, and so on.

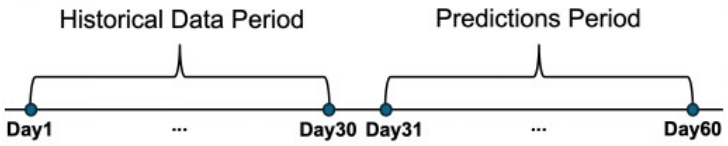

The Decision Task

You will make predictions on **whether the stock's closing price was up or down** for each trading day from Day 31 to Day 60. You will choose one option (either up or down) and if your choice matches the actual price change (i.e., up or down) compared to the previous trading day, you will win $2; otherwise, you will win nothing.

An example of what you will see in the main tasks: in the following, you will make the prediction on whether the stock's closing price was up or down on Day 31 compared to the previous trading day.

Please make the prediction on whether the stock's closing price was up or down on **Day 31** compared to the previous trading day.

Up Down

Suppose you choose Up, and the actual closing price was Up compared to the previous trading day, you will win $2.

Suppose you choose Up, and the actual closing price was Down compared to the previous trading day, you will win nothing.

Your Payoff

After completing the experiment, one of your predictions from Day 31 to Day 60 will be randomly drawn for the bonus payment. If your prediction for the drawn day is correct, you will earn **$2**; otherwise, you will receive no bonus.

Historical Data

The stock's closing price changes from Day 1 to Day 30 (compared to the previous trading day) are shown as follows:

Historical Data	
Day	**Closing Price Changes**
Day 1	Down
Day 2	Up
Day 3	Down
Day 4	Up
Day 5	Down
Day 6	Up
Day 7	Down
Day 8	Down
Day 9	Up
Day 10	Down
Day 11	Down
Day 12	Up
Day 13	Down
Day 14	Up
Day 15	Up
Day 16	Up
Day 17	Up
Day 18	Down
Day 19	Down
Day 20	Down
Day 21	Down
Day 22	Up
Day 23	Down
Day 24	Up
Day 25	Up
Day 26	Up
Day 27	Up
Day 28	Down
Day 29	Up
Day 30	Down

Historical data is available on each page where you make your predictions. Before the main tasks, we would like you to complete two practice tasks to ensure that you fully understand the rules. Please read the instructions above carefully and click the "Continue"

button to begin the practice section.

You will be able to continue after **30** seconds. Please read the instructions carefully.

(Continue)

(Page 3)

Practice Task 1

Historical Data

The stock's closing price changes from Day 1 to Day 30 (compared to the previous trading day) are shown as follows:

Historical Data	
Day	**Closing Price Changes**
Day 1	Down
Day 2	Up
Day 3	Down
Day 4	Up
Day 5	Down
Day 6	Up
Day 7	Down
Day 8	Down
Day 9	Up
Day 10	Down
Day 11	Down
Day 12	Up
Day 13	Down
Day 14	Up
Day 15	Up
Day 16	Up
Day 17	Up
Day 18	Down
Day 19	Down
Day 20	Down
Day 21	Down
Day 22	Up
Day 23	Down
Day 24	Up
Day 25	Up
Day 26	Up
Day 27	Up
Day 28	Down
Day 29	Up
Day 30	Down

Please make the prediction on whether the stock's closing price was up or down on **Day XX** compared to the previous trading day. Please choose one option.

Up Down

Note: you can click the following button to read the instructions for this task.

[Instructions]

(Page 4)

Practice Task 2

Historical Data

The stock's closing price changes from Day 1 to Day 30 (compared to the previous trading day) are shown as follows:

Historical Data	
Day	**Closing Price Changes**
Day 1	Down
Day 2	Up
Day 3	Down
Day 4	Up
Day 5	Down
Day 6	Up
Day 7	Down
Day 8	Down
Day 9	Up
Day 10	Down
Day 11	Down
Day 12	Up
Day 13	Down
Day 14	Up
Day 15	Up
Day 16	Up
Day 17	Up
Day 18	Down
Day 19	Down
Day 20	Down
Day 21	Down
Day 22	Up
Day 23	Down
Day 24	Up
Day 25	Up
Day 26	Up
Day 27	Up
Day 28	Down
Day 29	Up
Day 30	Down

Please make the prediction on whether the stock's closing price was up or down on **Day XX** compared to the previous trading day. Please choose one option.

Up Down

Note: you can click the following button to read the instructions for this task.

[Instructions]

(Page 5)

Practice Tasks Completed

You have completed the practice tasks. Please click the "Continue" button to begin the main tasks.

(Continue)

(Pages 6-35)[13]

Main Task - Day 31

Historical Data

The stock's closing price changes from Day 1 to Day 30 (compared to the previous trading day) are shown as follows:

Historical Data	
Day	**Closing Price Changes**
Day 1	Down
Day 2	Up
Day 3	Down
Day 4	Up
Day 5	Down
Day 6	Up
Day 7	Down
Day 8	Down
Day 9	Up
Day 10	Down
Day 11	Down
Day 12	Up
Day 13	Down
Day 14	Up
Day 15	Up
Day 16	Up
Day 17	Up
Day 18	Down
Day 19	Down
Day 20	Down
Day 21	Down
Day 22	Up
Day 23	Down
Day 24	Up
Day 25	Up
Day 26	Up
Day 27	Up
Day 28	Down
Day 29	Up
Day 30	Down

Please make the prediction on whether the stock's closing price was up or down on **Day 31** compared to the previous trading day. Please choose one option.

Up Down

Note: you can click the following button to read the instructions for this task.

[Instructions]

[13]Followed by the same separate pages as in the *Odd-Even* treatment.

C Experimental Instructions: Main Experiment

C.1 *Odd-Even* Treatment

(The First Page)[14]

Welcome to our Decision-Making Study!

The entire study, including tasks and a questionnaire, will take approximately **15** minutes to complete.

In the main tasks, you will participate in two stages of decision making: **Stage 1** and **Stage 2**. In each stage, you will make decisions for **5** rounds.

In addition to a participation fee, you have the opportunity to earn a bonus based on your decisions. After completing the experiment, one task involving payment will be randomly drawn, and your bonus will be determined by the decision you made in that task. If the drawn task includes multiple options, one option will be randomly drawn for payment.

Throughout the experiment, all amounts will be displayed in Experimental Currency Units (ECU). In the end, your bonus will be converted to USD at a rate of **1 USD = 10 ECU**. Your final payoff will include both your participation fee and the bonus from the randomly drawn task.

Please take your time to carefully read and thoughtfully respond to all the questions. All decisions and personal information in this study will remain completely anonymous and strictly confidential. Thank you for your participation!

Please wait 10 seconds...

(Continue)

[14]Each subject was presented with a consent form at the beginning of the experiment, and only those who accepted the consent form could proceed with the study. The setting is the same for other two treatments (i.e., *Up-Down* and *Numbers*).

(Stage 1 - The First Page)

Instructions of Stage 1

In this task, you will decide whether to follow the advice of an advisor, which can be either a **robo-advisor** or a **human advisor**, on whether the last digit of a stock in a randomly drawn date (between Day 31 and Day 60 as explained below) was odd or even. If you choose to follow an advisor's advice about the bet on odd or even, your payoff from the bet will depend on their advice.

The Stock

The stock is a real, publicly traded stock randomly selected from the stock exchange. This stock will remain the same throughout today's experiment. The 60 consecutive trading days are labeled as Day 1 to Day 60. The first trading day is Day 1, the second is Day 2, and so on.

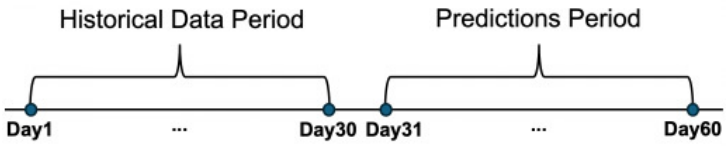

How Advisors Provide Advice

Robo-Advisor: We provided ChatGPT with historical data on the stock's closing prices for the trading days from Day 1 to Day 30, without disclosing the stock's ticker symbol. The ChatGPT provided 100 independent predictions on whether the last digit (i.e., the hundredths decimal place) of the stock's closing price was **odd** or **even** for each trading day from Day 31 to Day 60. For each prediction, the ChatGPT will choose one option (either odd or even) and if its choice matches the actual digit's parity (i.e., odd or even), the prediction is considered correct.

Human Advisor: We provided 100 real-life investors with historical data on the stock's closing prices for the trading days from Day 1 to Day 30, without disclosing the stock's ticker symbol. Each investor independently provided predictions on whether the last digit (i.e., the hundredths decimal place) of the stock's closing price was **odd** or **even** for each trading day from Day 31 to Day 60. For each prediction, the investor will choose one

option (either odd or even) and if her/his choice matches the actual digit's parity (i.e., odd or even), the prediction is considered correct.

The Decision Task

In this stage, you will make 5 rounds of decisions. In each round, a prediction trading day (between Day 31 and Day 60) will be randomly drawn to determine which trading day will be used for your bonus payment calculation. In each round, you will choose between two options:

1. Follow an advisor's advice.

If you choose to follow an advisor, one of the 100 predictions from the advisor on the drawn date will be randomly drawn as your prediction. If the prediction is correct for the randomly drawn date from Day 31 to Day 60, you will win the specified amount listed on the decision page; otherwise, you will receive the lower amount listed on the same page.

2. Receive a sure amount of money.

If you choose to receive the sure amount and this option is drawn as your bonus payment, you will win the specified amount listed on the decision page.

Note that the payoff for the bets will differ between rounds. Please pay attention to the payoff displayed on the decision page.

You will be able to continue after **30** seconds. Please read the instructions carefully.

(Continue)

(Stage 1 - Page 2[15])

Stage 1 - Round 1

[Step 1 / 3]

The task will proceed in three steps. Steps 1 and 2 will help you narrow down your choices. Step 3 will display your choices based on your choices in Steps 1 and 2. Please make your choices below. You can click the "Reset" button to undo your choices (those shown in blue on the screen). You can click the "Instructions" button at any time to review the instructions.

Instructions

Which option do you choose?			
Option 1 Follow the robo-advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 10 ECU if win, 0 ECU otherwise.	◉	○	0 ECU
	◉	○	5 ECU
	○	◉	10 ECU

Which option do you choose?			
Option 1 Follow the human advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 10 ECU if win, 0 ECU otherwise.	◉	○	0 ECU
	○	◉	5 ECU
	○	◉	10 ECU

Continue

Reset

[15]In the current page, we only present Round 1 out of 5 as an example. In each round, the payoff in Option 1 varies, with equal probability of randomly displaying one of the following five pairs: (10, 0), (15, 0), (25, 0), (20, 10), and (25, 15). The setting is the same for other two treatments (i.e., *Up-Down* and *Numbers*).

(Stage 1 - Page 3)

Stage 1 - Round 1

[Step 2 / 3]

The choices below (shown in grey) have been pre-filled based on your Step 1 choices. Please complete the remaining choices below, indicating choices between following the advisor's advice or receiving a specific sure amount. You can click the "Reset" button to undo your choices (those shown in blue on the screen). You may also click the "Back" button to revise your Step 1 choices. You can click the "Instructions" button at any time to review the instructions.

Instructions

Which option do you choose?			
Option 1 Follow the robo-advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
	◉	○	0 ECU
	◉	○	0.5 ECU
	◉	○	1 ECU
	◉	○	1.5 ECU
	◉	○	2 ECU
	◉	○	2.5 ECU
	◉	○	3 ECU
	◉	○	3.5 ECU
	◉	○	4 ECU
	◉	○	4.5 ECU
Receive 10 ECU if win, 0 ECU otherwise.	●	○	5 ECU
	○	●	5.5 ECU
	○	●	6 ECU
	○	●	6.5 ECU
	○	●	7 ECU
	○	●	7.5 ECU
	○	●	8 ECU
	○	●	8.5 ECU
	○	●	9 ECU
	○	●	9.5 ECU
	○	●	10 ECU

Which option do you choose?			
Option 1 Follow the human advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 10 ECU if win, 0 ECU otherwise.	◉	○	0 ECU
	○	◉	0.5 ECU
	○	◉	1 ECU
	○	◉	1.5 ECU
	○	◉	2 ECU
	○	◉	2.5 ECU
	○	◉	3 ECU
	○	◉	3.5 ECU
	○	◉	4 ECU
	○	◉	4.5 ECU
	○	◉	5 ECU
	○	◉	5.5 ECU
	○	◉	6 ECU
	○	◉	6.5 ECU
	○	◉	7 ECU
	○	◉	7.5 ECU
	○	◉	8 ECU
	○	◉	8.5 ECU
	○	◉	9 ECU
	○	◉	9.5 ECU
	○	◉	10 ECU

Continue

Reset

Back

(Stage 1 - Page 4)

Stage 1 - Round 1

[Step 3 / 3]

The choices below (shown in grey) are based on your choices in Steps 1 and 2. You can submit your final choices by clicking the "Confirm" button. You may also click the "Back" button to revise your Step 2 choices. You can click the "Instructions" button at any time to review the instructions.

Instructions

Which option do you choose?			
Option 1 Follow the robo-advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 10 ECU if win, 0 ECU otherwise.	◉	○	0 ECU
	◉	○	0.5 ECU
	◉	○	1 ECU
	◉	○	1.5 ECU
	◉	○	2 ECU
	◉	○	2.5 ECU
	◉	○	3 ECU
	◉	○	3.5 ECU
	◉	○	4 ECU
	◉	○	4.5 ECU
	◉	○	5 ECU
	○	◉	5.5 ECU
	○	◉	6 ECU
	○	◉	6.5 ECU
	○	◉	7 ECU
	○	◉	7.5 ECU
	○	◉	8 ECU
	○	◉	8.5 ECU
	○	◉	9 ECU
	○	◉	9.5 ECU
	○	◉	10 ECU

Which option do you choose?			
Option 1 Follow the human advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 10 ECU if win, 0 ECU otherwise.	◉	○	0 ECU
	○	◉	0.5 ECU
	○	◉	1 ECU
	○	◉	1.5 ECU
	○	◉	2 ECU
	○	◉	2.5 ECU
	○	◉	3 ECU
	○	◉	3.5 ECU
	○	◉	4 ECU
	○	◉	4.5 ECU
	○	◉	5 ECU
	○	◉	5.5 ECU
	○	◉	6 ECU
	○	◉	6.5 ECU
	○	◉	7 ECU
	○	◉	7.5 ECU
	○	◉	8 ECU
	○	◉	8.5 ECU
	○	◉	9 ECU
	○	◉	9.5 ECU
	○	◉	10 ECU

Confirm

Back

(Stage 2 - The First Page)

Instructions of Stage 2

In this stage, you will make 5 rounds of decisions. In each round, your previous decisions in Stage 1 on whether to follow an advisor's advice or receive a sure amount will be displayed. You will also be provided with information about the advisors' prediction accuracy, and you may revise your decisions if you want. Your bonus payment will be based on your revised choices in this stage. You can either update your decisions or keep the original ones displayed on the following pages.

The following instructions are the same as those in Stage 1:

In this task, you will decide whether to follow the advice of an advisor, which can be either a **robo-advisor** or a **human advisor**, on whether the last digit of a stock in a randomly drawn date (between Day 31 and Day 60 as explained below) was odd or even. If you choose to follow an advisor's advice about the bet on odd or even, your payoff from the bet will depend on their advice.

The Stock

The stock is a real, publicly traded stock randomly selected from the stock exchange. This stock will remain the same throughout today's experiment. The 60 consecutive trading days are labeled as Day 1 to Day 60. The first trading day is Day 1, the second is Day 2, and so on.

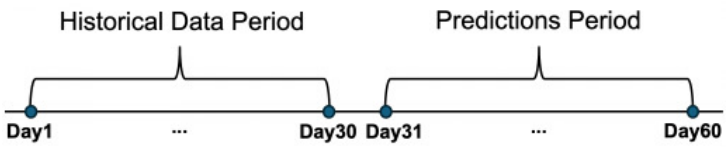

How Advisors Provide Advice

Robo-Advisor: We provided ChatGPT with historical data on the stock's closing prices for the trading days from Day 1 to Day 30, without disclosing the stock's ticker symbol. The ChatGPT provided 100 independent predictions on whether the last digit (i.e., the hundredths decimal place) of the stock's closing price was **odd** or **even** for each trading

day from Day 31 to Day 60. For each prediction, the ChatGPT will choose one option (either odd or even) and if its choice matches the actual digit's parity (i.e., odd or even), the prediction is considered correct.

Human Advisor: We provided 100 real-life investors with historical data on the stock's closing prices for the trading days from Day 1 to Day 30, without disclosing the stock's ticker symbol. Each investor independently provided predictions on whether the last digit (i.e., the hundredths decimal place) of the stock's closing price was **odd** or **even** for each trading day from Day 31 to Day 60. For each prediction, the investor will choose one option (either odd or even) and if her/his choice matches the actual digit's parity (i.e., odd or even), the prediction is considered correct.

The Decision Task

In this stage, you will make 5 rounds of decisions. In each round, a prediction trading day (between Day 31 and Day 60) will be randomly drawn to determine which trading day will be used for your bonus payment calculation. In each round, you will choose between two options:

1. Follow an advisor's advice.

If you choose to follow an advisor, one of the 100 predictions from the advisor on the drawn date will be randomly drawn as your prediction. If the prediction is correct for the randomly drawn date from Day 31 to Day 60, you will win the specified amount listed on the decision page; otherwise, you will receive the lower amount listed on the same page.

2. Receive a sure amount of money.

If you choose to receive the sure amount and this option is drawn as your bonus payment, you will win the specified amount listed on the decision page.

Note that the payoff for the bets will differ between rounds. Please pay attention to the payoff displayed on the decision page.

You will be able to continue after **30** seconds. Please read the instructions carefully.

(Continue)

(Stage 2 - Page 2[16])

Stage 2 - Round 1

We would like to provide you with information about the advisors' prediction accuracy:

Based on the data, on average, **49.67%** of the **robo-advisor**'s predictions are correct on any given trading day.

Based on the data, on average, **49.77%** of the **human advisor**'s predictions are correct on any given trading day.

Your previous decisions in Stage 1 on whether to follow an advisor's advice or receive a sure amount are displayed on this page. You can either update your decisions or keep the original ones. Once you have finished your revisions or decided to keep your previous decisions, please click the "Submit" button. If you want to undo your changes and restore your original choices, you may click the "Reset to Previous Decisions" button.

Note: you can click the following button to read the instructions for this task.

Instructions

Which option do you choose?			
Option 1 Follow the human advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 10 ECU if win, 0 ECU otherwise.	◉	○	0 ECU
	○	◉	0.5 ECU
	○	◉	1 ECU
	○	◉	1.5 ECU
	○	◉	2 ECU
	○	◉	2.5 ECU
	○	◉	3 ECU
	○	◉	3.5 ECU
	○	◉	4 ECU
	○	◉	4.5 ECU
	○	◉	5 ECU
	○	◉	5.5 ECU
	○	◉	6 ECU
	○	◉	6.5 ECU
	○	◉	7 ECU
	○	◉	7.5 ECU
	○	◉	8 ECU
	○	◉	8.5 ECU
	○	◉	9 ECU
	○	◉	9.5 ECU
	○	◉	10 ECU

[16]In the current page, we only present Round 1 out of 5 as an example. In Stage 2, the order of payoffs follows exactly the same sequence as in Stage 1. However, the order of the Human Advisor block and the Robo-Advisor block is still randomized. The setting is the same for other two treatments (i.e., *Up-Down* and *Numbers*).

Which option do you choose?			
Option 1 Follow the robo-advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 10 ECU if win, 0 ECU otherwise.	◉	○	0 ECU
	◉	○	0.5 ECU
	◉	○	1 ECU
	◉	○	1.5 ECU
	◉	○	2 ECU
	◉	○	2.5 ECU
	◉	○	3 ECU
	◉	○	3.5 ECU
	◉	○	4 ECU
	◉	○	4.5 ECU
	◉	○	5 ECU
	○	◉	5.5 ECU
	○	◉	6 ECU
	○	◉	6.5 ECU
	○	◉	7 ECU
	○	◉	7.5 ECU
	○	◉	8 ECU
	○	◉	8.5 ECU
	○	◉	9 ECU
	○	◉	9.5 ECU
	○	◉	10 ECU

Submit

Reset to Previous Decisions

(Other Tasks - The First Page)

How well do you know the following technologies before the experiment?

	Very unfamiliar	Somewhat unfamiliar	Neither familiar nor unfamiliar	Somewhat familiar	Very familiar
Artificial Intelligence	○ 1	○ 2	○ 3	○ 4	○ 5
Blockchain	○ 1	○ 2	○ 3	○ 4	○ 5
Cloud Computing	○ 1	○ 2	○ 3	○ 4	○ 5
Big Data	○ 1	○ 2	○ 3	○ 4	○ 5
Digital Currency	○ 1	○ 2	○ 3	○ 4	○ 5

(Continue)

(Other Tasks - Page 2)

This is an attention question. Please choose the option that contains the number three.

❍ 6

❍ 3

❍ 7

❍ 4

(Continue)

(Other Tasks - Page 3a)

Are you familiar with the robo-advisor?(1: very unfamiliar, ..., 10: very familiar)

❍ 1 ❍ 2 ❍ 3 ❍ 4 ❍ 5 ❍ 6 ❍ 7 ❍ 8 ❍ 9 ❍ 10

(Continue)

(Other Tasks - Page 3b)

Are you familiar with the human advisor?(1: very unfamiliar, ..., 10: very familiar)

❍ 1 ❍ 2 ❍ 3 ❍ 4 ❍ 5 ❍ 6 ❍ 7 ❍ 8 ❍ 9 ❍ 10

(Continue)

(Other Tasks - Page 4a)

Do you have any prior experience using the robo-advisor?

❍ Yes

❍ No

(Continue)

(Other Tasks - Page 4b)

Do you have any prior experience using the human advisor?

❍ Yes

❍ No

(Continue)

(Other Tasks - Page 5a)

Whether your prior experience with the robo-advisor is good or bad?(1: very bad, ..., 10: very good)

❍ 1 ❍ 2 ❍ 3 ❍ 4 ❍ 5 ❍ 6 ❍ 7 ❍ 8 ❍ 9 ❍ 10

(Continue)

(Other Tasks - Page 5b)

Whether your prior experience with the human advisor is good or bad?(1: very bad, ..., 10: very good)

❍ 1 ❍ 2 ❍ 3 ❍ 4 ❍ 5 ❍ 6 ❍ 7 ❍ 8 ❍ 9 ❍ 10

(Continue)

(Questionnaire - The First Page)

This section includes a questionnaire that we would like you to complete. Please answer the questions thoughtfully and accurately. Your responses will be kept confidential and anonymous. Thank you!

(Continue)

(Questionnaire - Page 2)

Questionnaire

What is your gender?

○ Male

○ Female

(Continue)

(Questionnaire - Page 3)

Questionnaire

What is your age?

_____

(Continue)

(Questionnaire - Page 4)

Questionnaire

What is your annual income?

○ Less than $10,000

○ Between $10,000 - $19,999

○ Between $20,000 - $29,999

○ Between $30,000 - $39,999

○ Between $40,000 - $49,999

○ Between $50,000 - $59,999

○ Between $60,000 - $69,999

○ Between $70,000 - $79,999

○ Between $80,000 - $89,999

○ Between $90,000 - $99,999

○ Between $100,000 - $149,999

○ More than $150,000

(Continue)

(Questionnaire - Page 5)

Questionnaire

What is the highest degree or level of school you have completed?

○ Some elementary school or below

○ Elementary school

○ Middle school

○ High school or equivalent

○ Some college/associate degree

○ Bachelor's degree

○ Master's degree

○ Professional or PhD

(Continue)

(Questionnaire - Page 6)

Questionnaire

Are you currently...?

○ Employed for wages

○ Self-employed

○ Out of work and looking for work

○ Out of work but not currently looking for work

○ A homemaker

○ A student

○ Military

○ Retired

○ Unable to work

(Continue)

(The Last Page)

Experiment Finished

Thank you for participating in this experiment!

Please do not forget to click the "Continue and confirm your participation" button in the lower right corner of this page. Only in this way will your submission be successfully received, and only the successful submission will be approved and paid.

(Continue and confirm your participation)

C.2 *Up-Down* Treatment

(The First Page)

Welcome to our Decision-Making Study!

The entire study, including tasks and a questionnaire, will take approximately **15** minutes to complete.

In the main tasks, you will participate in two stages of decision making: **Stage 1** and **Stage 2**. In each stage, you will make decisions for **5** rounds.

In addition to a participation fee, you have the opportunity to earn a bonus based on your decisions. After completing the experiment, one task involving payment will be randomly drawn, and your bonus will be determined by the decision you made in that task. If the drawn task includes multiple options, one option will be randomly drawn for payment.

Throughout the experiment, all amounts will be displayed in Experimental Currency Units (ECU). In the end, your bonus will be converted to USD at a rate of **1 USD = 10 ECU**. Your final payoff will include both your participation fee and the bonus from the randomly drawn task.

Please take your time to carefully read and thoughtfully respond to all the questions. All decisions and personal information in this study will remain completely anonymous and strictly confidential. Thank you for your participation!

Please wait 10 seconds...

(Continue)

(Stage 1 - The First Page)

Instructions of Stage 1

In this task, you will decide whether to follow the advice of an advisor, which can be either a **robo-advisor** or a **human advisor**, on whether a stock's closing price in a randomly drawn date (between Day 31 and Day 60 as explained below) was up or down compared to the previous trading day. If you choose to follow an advisor's advice about the bet on up or down, your payoff from the bet will depend on their advice.

The Stock

The stock is a real, publicly traded stock randomly selected from the stock exchange. This stock will remain the same throughout today's experiment. The 60 consecutive trading days are labeled as Day 1 to Day 60. The first trading day is Day 1, the second is Day 2, and so on.

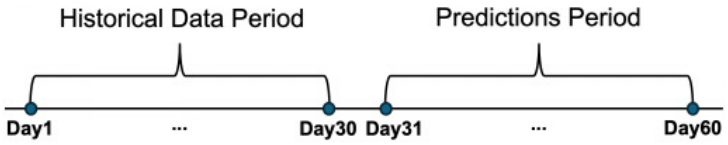

How Advisors Provide Advice

Robo-Advisor: We provided ChatGPT with historical data on the stock's closing price changes for the trading days from Day 1 to Day 30 compared to the previous trading day, without disclosing the stock's ticker symbol. The ChatGPT provided 100 independent predictions on whether the stock's closing price was **up** or **down** for each trading day from Day 31 to Day 60 compared to the previous trading day. For each prediction, the ChatGPT will choose one option (either up or down) and if its choice matches the actual price change (i.e., up or down) compared to the previous trading day, the prediction is considered correct.

Human Advisor: We provided 100 real-life investors with historical data on the stock's closing price changes for the trading days from Day 1 to Day 30 compared to the previous trading day, without disclosing the stock's ticker symbol. Each investor independently provided predictions on whether the stock's closing price was **up** or **down** for each trading day from Day 31 to Day 60 compared to the previous trading day. For each prediction,

the investor will choose one option (either up or down) and if her/his choice matches the actual price change (i.e., up or down) compared to the previous trading day, the prediction is considered correct.

The Decision Task

In this stage, you will make 5 rounds of decisions. In each round, a prediction trading day (between Day 31 and Day 60) will be randomly drawn to determine which trading day will be used for your bonus payment calculation. In each round, you will choose between two options:

1. Follow an advisor's advice.

If you choose to follow an advisor, one of the 100 predictions from the advisor on the drawn date will be randomly drawn as your prediction. If the prediction is correct for the randomly drawn date from Day 31 to Day 60, you will win the specified amount listed on the decision page; otherwise, you will receive the lower amount listed on the same page.

2. Receive a sure amount of money.

If you choose to receive the sure amount and this option is drawn as your bonus payment, you will win the specified amount listed on the decision page.

Note that the payoff for the bets will differ between rounds. Please pay attention to the payoff displayed on the decision page.

You will be able to continue after **30** seconds. Please read the instructions carefully.

(Continue)

(Stage 1 - Page 2)

Stage 1 - Round 1

[Step 1 / 3]

The task will proceed in three steps. Steps 1 and 2 will help you narrow down your choices. Step 3 will display your choices based on your choices in Steps 1 and 2. Please make your choices below. You can click the "Reset" button to undo your choices (those shown in blue on the screen). You can click the "Instructions" button at any time to review the instructions.

Instructions

Which option do you choose?			
Option 1 Follow the robo-advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 25 ECU if win, 15 ECU otherwise.	◉	○	15 ECU
	◉	○	20 ECU
	○	◉	25 ECU

Which option do you choose?			
Option 1 Follow the human advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 25 ECU if win, 15 ECU otherwise.	◉	○	15 ECU
	○	◉	20 ECU
	○	◉	25 ECU

Continue

Reset

(Stage 1 - Page 3)

Stage 1 - Round 1

[Step 2 / 3]

The choices below (shown in grey) have been pre-filled based on your Step 1 choices. Please complete the remaining choices below, indicating choices between following the advisor's advice or receiving a specific sure amount. You can click the "Reset" button to undo your choices (those shown in blue on the screen). You may also click the "Back" button to revise your Step 1 choices. You can click the "Instructions" button at any time to review the instructions.

Instructions

Which option do you choose?			
Option 1 Follow the robo-advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
	◉	○	15 ECU
	◉	○	15.5 ECU
	◉	○	16 ECU
	◉	○	16.5 ECU
	◉	○	17 ECU
	◉	○	17.5 ECU
	◉	○	18 ECU
	◉	○	18.5 ECU
	◉	○	19 ECU
	◉	○	19.5 ECU
Receive 25 ECU if win, 15 ECU otherwise.	●	○	20 ECU
	○	●	20.5 ECU
	○	●	21 ECU
	○	●	21.5 ECU
	○	●	22 ECU
	○	●	22.5 ECU
	○	●	23 ECU
	○	●	23.5 ECU
	○	●	24 ECU
	○	●	24.5 ECU
	○	●	25 ECU

Which option do you choose?			
Option 1 Follow the human advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 25 ECU if win, 15 ECU otherwise.	○	●	15 ECU
	○	●	15.5 ECU
	○	●	16 ECU
	○	●	16.5 ECU
	○	●	17 ECU
	○	●	17.5 ECU
	○	●	18 ECU
	○	●	18.5 ECU
	○	●	19 ECU
	○	●	19.5 ECU
	○	●	20 ECU
	○	●	20.5 ECU
	○	●	21 ECU
	○	●	21.5 ECU
	○	●	22 ECU
	○	●	22.5 ECU
	○	●	23 ECU
	○	●	23.5 ECU
	○	●	24 ECU
	○	●	24.5 ECU
	○	●	25 ECU

Continue

Reset

Back

(Stage 1 - Page 4)

Stage 1 - Round 1

[Step 3 / 3]

The choices below (shown in grey) are based on your choices in Steps 1 and 2. You can submit your final choices by clicking the "Confirm" button. You may also click the "Back" button to revise your Step 2 choices. You can click the "Instructions" button at any time to review the instructions.

Instructions

Which option do you choose?			
Option 1 Follow the robo-advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 25 ECU if win, 15 ECU otherwise.	◉	○	15 ECU
	◉	○	15.5 ECU
	◉	○	16 ECU
	◉	○	16.5 ECU
	◉	○	17 ECU
	◉	○	17.5 ECU
	◉	○	18 ECU
	◉	○	18.5 ECU
	◉	○	19 ECU
	◉	○	19.5 ECU
	◉	○	20 ECU
	○	◉	20.5 ECU
	○	◉	21 ECU
	○	◉	21.5 ECU
	○	◉	22 ECU
	○	◉	22.5 ECU
	○	◉	23 ECU
	○	◉	23.5 ECU
	○	◉	24 ECU
	○	◉	24.5 ECU
	○	◉	25 ECU

Which option do you choose?			
Option 1 Follow the human advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 25 ECU if win, 15 ECU otherwise.	◉	○	15 ECU
	○	◉	15.5 ECU
	○	◉	16 ECU
	○	◉	16.5 ECU
	○	◉	17 ECU
	○	◉	17.5 ECU
	○	◉	18 ECU
	○	◉	18.5 ECU
	○	◉	19 ECU
	○	◉	19.5 ECU
	○	◉	20 ECU
	○	◉	20.5 ECU
	○	◉	21 ECU
	○	◉	21.5 ECU
	○	◉	22 ECU
	○	◉	22.5 ECU
	○	◉	23 ECU
	○	◉	23.5 ECU
	○	◉	24 ECU
	○	◉	24.5 ECU
	○	◉	25 ECU

Confirm

Back

(Stage 2 - The First Page)

Instructions of Stage 2

In this stage, you will make 5 rounds of decisions. In each round, your previous decisions in Stage 1 on whether to follow an advisor's advice or receive a sure amount will be displayed. You will also be provided with information about the advisors' prediction accuracy, and you may revise your decisions if you want. Your bonus payment will be based on your revised choices in this stage. You can either update your decisions or keep the original ones displayed on the following pages.

The following instructions are the same as those in Stage 1:

In this task, you will decide whether to follow the advice of an advisor, which can be either a **robo-advisor** or a **human advisor**, on whether a stock's closing price in a randomly drawn date (between Day 31 and Day 60 as explained below) was up or down compared to the previous trading day. If you choose to follow an advisor's advice about the bet on up or down, your payoff from the bet will depend on their advice.

The Stock

The stock is a real, publicly traded stock randomly selected from the stock exchange. This stock will remain the same throughout today's experiment. The 60 consecutive trading days are labeled as Day 1 to Day 60. The first trading day is Day 1, the second is Day 2, and so on.

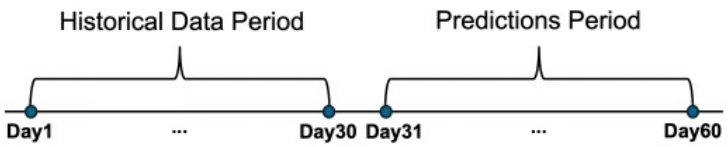

How Advisors Provide Advice

Robo-Advisor: We provided ChatGPT with historical data on the stock's closing price changes for the trading days from Day 1 to Day 30 compared to the previous trading day, without disclosing the stock's ticker symbol. The ChatGPT provided 100 independent predictions on whether the stock's closing price was **up** or **down** for each trading day

from Day 31 to Day 60 compared to the previous trading day. For each prediction, the ChatGPT will choose one option (either up or down) and if its choice matches the actual price change (i.e., up or down) compared to the previous trading day, the prediction is considered correct.

Human Advisor: We provided 100 real-life investors with historical data on the stock's closing price changes for the trading days from Day 1 to Day 30 compared to the previous trading day, without disclosing the stock's ticker symbol. Each investor independently provided predictions on whether the stock's closing price was **up** or **down** for each trading day from Day 31 to Day 60 compared to the previous trading day. For each prediction, the investor will choose one option (either up or down) and if her/his choice matches the actual price change (i.e., up or down) compared to the previous trading day, the prediction is considered correct.

The Decision Task

In this stage, you will make 5 rounds of decisions. In each round, a prediction trading day (between Day 31 and Day 60) will be randomly drawn to determine which trading day will be used for your bonus payment calculation. In each round, you will choose between two options:

1. Follow an advisor's advice.

If you choose to follow an advisor, one of the 100 predictions from the advisor on the drawn date will be randomly drawn as your prediction. If the prediction is correct for the randomly drawn date from Day 31 to Day 60, you will win the specified amount listed on the decision page; otherwise, you will receive the lower amount listed on the same page.

2. Receive a sure amount of money.

If you choose to receive the sure amount and this option is drawn as your bonus payment, you will win the specified amount listed on the decision page.

Note that the payoff for the bets will differ between rounds. Please pay attention to the payoff displayed on the decision page.

You will be able to continue after **30** seconds. Please read the instructions carefully.

(Continue)

(Stage 2 - Page 2)

Stage 2 - Round 1

We would like to provide you with information about the advisors' prediction accuracy:

Based on the data, on average, **50.20%** of the **robo-advisor**'s predictions are correct on any given trading day.

Based on the data, on average, **54.20%** of the **human advisor**'s predictions are correct on any given trading day.

Your previous decisions in Stage 1 on whether to follow an advisor's advice or receive a sure amount are displayed on this page. You can either update your decisions or keep the original ones. Once you have finished your revisions or decided to keep your previous decisions, please click the "Submit" button. If you want to undo your changes and restore your original choices, you may click the "Reset to Previous Decisions" button.

Note: you can click the following button to read the instructions for this task.

Instructions

Which option do you choose?			
Option 1 Follow the human advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 25 ECU if win, 15 ECU otherwise.	◉	○	15 ECU
	○	◉	15.5 ECU
	○	◉	16 ECU
	○	◉	16.5 ECU
	○	◉	17 ECU
	○	◉	17.5 ECU
	○	◉	18 ECU
	○	◉	18.5 ECU
	○	◉	19 ECU
	○	◉	19.5 ECU
	○	◉	20 ECU
	○	◉	20.5 ECU
	○	◉	21 ECU
	○	◉	21.5 ECU
	○	◉	22 ECU
	○	◉	22.5 ECU
	○	◉	23 ECU
	○	◉	23.5 ECU
	○	◉	24 ECU
	○	◉	24.5 ECU
	○	◉	25 ECU

Which option do you choose?			
Option 1 Follow the robo-advisor's advice	**1**	**2**	**Option 2** Receive this amount for sure
Receive 25 ECU if win, 15 ECU otherwise.	◉	○	15 ECU
	◉	○	15.5 ECU
	◉	○	16 ECU
	◉	○	16.5 ECU
	◉	○	17 ECU
	◉	○	17.5 ECU
	◉	○	18 ECU
	◉	○	18.5 ECU
	◉	○	19 ECU
	◉	○	19.5 ECU
	◉	○	20 ECU
	○	◉	20.5 ECU
	○	◉	21 ECU
	○	◉	21.5 ECU
	○	◉	22 ECU
	○	◉	22.5 ECU
	○	◉	23 ECU
	○	◉	23.5 ECU
	○	◉	24 ECU
	○	◉	24.5 ECU
	○	◉	25 ECU

Submit

Reset to Previous Decisions