FinTech Brings a Bias from Psychology Labs to a Two-trillion-dollar Market*

Hongjun Yan
DePaul University

Zhengwei Wang
Tsinghua University

Yulai Yuan
Xuanyuan Insurance Agency Co., Ltd

Xiaole Qiu
Tsinghua University

Jia Xiang
Tsinghua University

Yaqing Xiao
Capital University of Economics and Business

Erik Davidson
Baylor University

This Draft: May 26, 2026

*We are grateful for the helpful comments from James Choi, Lawrence Jin, Erik Mayer, Clemens Otto, Lin Peng, Alberto Rossi, Giulio Trigilia, Baolian Wang, Michael Weber, Edward Van Wesep, Arthur Yan, Dexin Zhou, Harald Uhlig, Bohui Zhang, and seminar participants at Baruch College, Cornell University, George Mason University, National University of Singapore, Singapore Management University, University of Chicago, University of Hong Kong, University of Illinois Chicago, University of Rochester, University of Wisconsin Madison, ABFER 13th Annual Conference, and SFS Cavalcade North America. We thank Lili Huang, Chunyan Wang, Jun Xiao, Liang Zhang, and an anonymous robo-advisory company for the collaborations on the RCT in this study and Charles Rotblut and the AAII for the collaboration on the survey. The RCT in this study was pre-registered with AEA RCT Registry (RCT ID: AEARCTR-0014422) and received IRB approval from the DePaul University (Protocol ID: IRB-2024-1411). Yan thanks the Fintech Dauphine Chair in partnership with Forvis Mazars and Crédit Agricole CIB. Please send all correspondence to Hongjun Yan, hongjun.yan.2011@gmail.com.

FinTech Brings a Bias from Psychology Labs to a Two-trillion-dollar Market

Abstract

Personalized services have traditionally been difficult to automate and therefore costly. Modern technology, such as robo-advising, promises to change this but also introduces new challenges. Robo-advisors typically make investment recommendations based on surveys of clients' preferences. As a result, biases induced by these surveys can become embedded in the algorithmic recommendations and influence investors' financial decisions. To examine this conjecture, we conduct two studies. First, we administer a nationwide survey of experienced U.S. investors and show that the response-order effect—the influence of the order of listed choices on respondents' answers—applies to survey questions commonly used by robo-advisors. Second, in collaboration with a leading robo-advisor in China, we conduct a preregistered RCT and show that, despite the large financial stakes involved, the response-order effect significantly alters both the robo-advisor's risk assessments and its clients' actual investment decisions, raising new ethical concerns in the FinTech era.

JEL Classification Numbers: D1, D9, G10, G2, G23, G40, G41.

Keywords: FinTech, Robo-advisor, Order effect, RCT, Nudge, Survey, Interface.

1. Introduction

Personalized services have long been labor-intensive and resistant to automation. The Baumol effect suggests that such sectors tend to exhibit limited productivity growth and experience rising relative costs over time (Baumol, 1967). Indeed, personalized services have traditionally been affordable only to wealthy individuals. Modern technology, however, promises to change that. In recent years, numerous firms have started offering large-scale, low-cost, personalized services across a wide range of domains, including financial investments, insurance, retirement planning, housing decisions, career guidance, and even personalized nutrition.

While this development has transformed the service industry, it may also create new challenges. To deliver low-cost personalized services, providers cannot afford extensive personal interactions with each customer. Instead, they typically rely on brief surveys to collect essential information. Consequently, any biases induced by these surveys are embedded in the service itself and may ultimately shape customers' decisions.

Our paper provides one such example. Financial technology (FinTech) has fueled the rapid expansion of the robo-advisory market, which manages over $2 trillion in assets as of 2025. This innovation has the potential to democratize financial advice by delivering personalized investment recommendations at scale to households with modest wealth (D'Acunto and Rossi, 2023). However, our analysis shows that it may also introduce an innocuous behavioral bias—long documented in psychology laboratories—into this two-trillion-dollar market.

We focus on the well-known *response-order effect*: the order in which response options appear in a survey question can influence how respondents answer (Krosnick and Presser, 2010). We conjecture that this bias also operates in robo-advisory surveys, despite the large financial stakes involved. For instance, in a question assessing risk tolerance, presenting the most risk-tolerant option at the top rather than at the bottom of the list may elicit different responses from the same individual. Because robo-advisors use these survey responses as inputs to their algorithms, they may—perhaps unknowingly—embed the bias into their investment recommendations to clients. More importantly, we conjecture that this bias is not confined to the virtual world but extends to clients' real-world investment decisions.

We examine these conjectures in two studies. First, we administer a nationwide survey of experienced U.S. investors and document a substantial response-order effect in the risk-tolerance

questions commonly used by robo-advisors. Second, motivated by this survey evidence, we collaborate with a leading robo-advisor in China to conduct a preregistered randomized controlled trial (RCT) that provides causal evidence that response order significantly affects both the robo-advisor's risk assessments and its clients' investment decisions.

Our first study examines whether the response-order effect documented in psychology also arises in a robo-advisory setting. Two features distinguish our design from typical psychology studies: it centers on the survey questions commonly used by robo-advisors, and it recruits experienced investors rather than students or the general public. Specifically, we administered a survey of over 3,000 experienced investors across the U.S. We implemented two otherwise identical versions of the questionnaire that differed only in the order of response options. Each respondent was randomly assigned, with equal probability, to one of the two versions.

We find a strong response-order effect: respondents tend to select options near the top of the lists, leading to a systematic bias in survey results. For example, consider the following risk tolerance question, which is often used by robo-advisors: "Recognizing that investments can fluctuate in value, for a $100,000 investment, how much of a decline in the value of your investment portfolio could you tolerate over the course of a year?" The response options were presented in the following two versions:

Original-order survey	*Reverse-order survey*
−10% (−$10,000)	−40% (−$40,000)
−20% (−$20,000)	−30% (−$30,000)
−30% (−$30,000)	−20% (−$20,000)
−40% (−$40,000)	−10% (−$10,000)

Because respondents tend to select options near the top of the list, those shown the reverse-order survey should, on average, report higher risk tolerance. Indeed, 39% of respondents to the reverse-order survey (611 out of 1,576) selected the most risk tolerant answer—"−40% (−$40,000)". In contrast, when that option appeared at the bottom of the list in the original-order survey, only 28% of respondents (439 out of 1,550) selected it. This difference is highly statistically significant ($p = 6.29 \times 10^{-10}$). More broadly, we find a quantitatively large response-order effect across a variety of survey questions, including those on investment preferences and expectations. By contrast, the response-order effect disappears for objective questions with clear answers, such as demographics.

In the second study, we collaborate with a leading robo-advisor in China (which will be referred to as "the Advisor" hereafter) to conduct an RCT to examine whether the response-order

effect persists when large financial stakes are involved, and, more importantly, whether it has a meaningful effect on investors' financial decisions.

When prospective clients open accounts with the Advisor, they must complete an 11-question survey covering investment experience, risk tolerance, and related topics. Existing clients are also prompted to retake the survey periodically, typically once every two years. For each question, the client receives a numerical score based on his response. The total score across all 11 questions determines the client's *risk score*, which is then used to classify clients into five *risk categories*—C1 through C5—with C5 indicating the highest risk tolerance.

During the trial, the Advisor administered two versions of the survey that differed only in the order of response options for each question. In the original-order version, the option corresponding to the lowest score appeared at the top of the list, whereas in the reverse-order version, the option corresponding to the highest score appeared at the top. Each client was randomly assigned, with equal probability, to one of the two survey versions. Randomization was implemented based on clients' account numbers: those with odd numbers received the reverse-order survey, and those with even numbers received the original-order version.

New clients received the treatment between October 25 and December 27, 2024. For existing clients—those who opened accounts before October 25, 2024—the treatment was implemented in two waves: on November 1, 2024, 5,000 users were randomly selected and invited to retake the risk assessment survey, and the same procedure was repeated on December 18, 2024 with another 5,000 users. The trial concluded on April 1, 2025, with 10,283 users in total. Of these, 7,141 (5,818 existing users and 1,323 new users) received the treatment and form our *RCT sample*. The original-order and reverse-order users in this sample are well balanced in both sample size and observable characteristics. The remaining 3,142 users were not treated; they opened accounts between December 28, 2024 and April 1, 2025, when the odd–even account-number rule no longer determined survey assignment. We refer to this group as the *placebo sample*.

To estimate the order effect, we regress risk scores on $Reverse$, a dummy variable equal to one if the account number is odd (and the user therefore received the reverse-order survey) and zero otherwise. In the RCT sample, the estimated coefficient on $Reverse$ is 3.35 (s.e. = 0.31). That is, users assigned to the reverse-order survey have risk scores that are, on average, 3.35 points higher than those assigned to the original-order survey. As a falsification test, we perform the same analysis using the placebo sample, where the odd–even account-number rule no longer determines

survey assignment and therefore should have no predictive power. Consistent with the experimental design, we find no difference in risk scores between users with odd account numbers ($Reverse = 1$) and those with even account numbers ($Reverse = 0$).

Since a user's risk category is mechanically determined by the risk score, the estimated order effect on risk scores naturally translates into shifts in users' risk categories. For example, the share of users in the most risk-tolerant category, C5, rises from 47.9% (1,697 users) in the original-order group to 56.5% (2,035 users) in the reverse-order group—an increase of 8.6 percentage points. The difference is highly statistically significant ($p = 2.83 \times 10^{-13}$). Thus, despite the large financial stakes involved, a seemingly minor design feature—the order in which survey responses are presented—can have a substantial impact on users' reported risk preferences.

These results raise a natural question: does response order also affect users' actual investment decisions? Two mechanisms suggest that it may. First (the *nudging channel*), because risk assessments are an input to the Advisor's recommendation algorithm, the order effect becomes embedded in those recommendations, potentially nudging users toward different investment choices. Second (the *compliance channel*), the Advisor assigns users to risk categories (C1–C5) and classifies products by risk levels (R1–R5). For regulatory reasons, users may purchase only products with ratings at or below their assigned risk category. Thus, if response order shifts users into lower risk categories, it restricts the set of products they can purchase. As a result, the response-order effect on risk categories may carry over to users' investment decisions.

To analyze the order effect on investment decisions, we focus on two primary types of products at the Advisor: robo-advised and self-directed. The Advisor provides three robo-advised investment strategies: "Aggressive," "Balanced," and "Stable" strategies, which have high, medium, and low risk levels, respectively. Each strategy aims to build and maintain a target portfolio tailored to a specific risk level. For example, the target allocation to equities is 47% under the Aggressive strategy but only 24% under the Stable strategy.

To invest in a strategy, a user needs to set up a "special purpose account" dedicated to the strategy. A user can choose any one or multiple strategies. However, his risk category may steer him toward a particular strategy. For example, a C5 user's assessment report begins with: "Your risk profile is Aggressive" before explaining how Aggressive investors typically invest. The phrase for the risk category—Aggressive—is the same as that for the strategy. Moreover, if a C4 user attempts to set up an account for the Aggressive strategy, he would be reminded that the strategy's

risk level exceeds his category. Although the user can ignore this warning to set up the account, compliance restrictions will still limit the user's investment in the strategy: Whenever the strategy's recommendation includes purchasing an R5 product, the user would be reminded that he can only make the purchase after retaking the survey to obtain a C5 category. This impediment may limit the user's investment in the aggressive strategy.

The second type of investment is self-directed. It includes mutual funds, money market funds, and wealth management products offered by third parties, with the Advisor acting as an intermediary. In this case, users must actively make all decisions related to selecting and trading these products. In contrast, robo-advised investments are semi-automatic: users can accept or reject the Advisor's recommended transactions but cannot initiate alternatives. Because self-directed investing requires more active decision-making, the response-order effect is expected to be weaker for these investments than for robo-advised ones.

We first examine the order effect on robo-advised investments. To estimate the effect along the intensive margin, we focus on users who purchased robo-advised products during the trial, comprising 1,746 users with total purchases of RMB 48.59 million.[1] Regressing a user's total investment in the Aggressive strategy on $Reverse$, we find a coefficient of 3.28 (s.e. = 1.66), indicating that exposure to the reverse-order survey increases investment in the Aggressive strategy by approximately RMB 3,280 on average. This effect is economically meaningful relative to the average purchase size of RMB 27,840.

Our hypothesis is that this effect operates through response-order-induced variation in users' risk categories. To examine this mechanism, we use survey response order as an instrument for risk classification and decompose it into two components: an exogenous component driven by response order and a residual component reflecting users' underlying risk tolerance. Consistent with this hypothesis, both components positively predict investment in the Aggressive strategy. In particular, a one-percentage-point increase in the probability of being classified into the C5 category due to the exogenous manipulation increases purchases of the Aggressive product by RMB 443 on average.

We next examine the extensive margin. Using the full RCT sample, we construct a dummy variable $DRobo_i$ equal to one if user i invests in any robo-advised product during the trial and zero

[1] The exchange rate was USD 1 = RMB 7.27 on April 1, 2025.

otherwise, and regress it on $Reverse$. The estimated coefficient is −0.0254 (s.e. = 0.0105), indicating that users assigned to the reverse-order survey are 2.54 percentage points less likely to invest. This pattern is unlikely to be driven by the compliance channel, as reverse-order users tend to receive higher risk categories and thus face fewer investment restrictions. A more plausible explanation is that nudging backfires: inflated risk classifications steer users toward strategies that some perceive as too risky, leading them to opt out of investing altogether. Consistent with this interpretation, participation increases with underlying risk tolerance but decreases with the exogenous, survey-induced component of risk classification. Together, these results show that response-order-induced distortions shift both margins: they increase risk-taking conditional on participation while reducing overall participation.

We also conduct a parallel analysis for self-directed investments and find that all estimated effects have the same sign as their counterparts for robo-advised investments but are weaker in both statistical and economic terms. A likely explanation is that users play a more active role in decision-making for self-directed investments, which attenuates the response-order effect. This intuition also suggests that the response-order effect is likely to be even stronger in the U.S. robo-advisory market. In the U.S. setting, investors typically play a more passive role in the decision-making process. Once the risk-assessment survey is completed, investment decisions are executed automatically based on the survey results (Reher and Sokolinski, 2024). Consequently, the order effect on risk assessments is transmitted directly to actual investments.

We note that, in principle, similar effects could arise even with a human financial advisor. If an advisor uses a survey to elicit investors' goals and preferences and relies on these responses when forming recommendations, response-order effects may influence investment decisions in much the same way. FinTech is therefore not a necessary condition for such effects to influence investment outcomes. However, it plays a critical role by standardizing the decision-making process, thereby allowing these effects to operate systematically and at scale.

Our findings raise an important ethical concern. Because survey-elicited preferences are susceptible to manipulation, a robo-advisor can influence investors' purchases while preserving the appearance that these choices reflect investors' own preferences. Hence, a robo-advisor may profit by steering its clients toward products that are suboptimal from the clients' perspective. Moreover, our evidence indicates that preference constructs such as risk aversion are not fully determinate: their estimated values—whether inferred from survey responses or investment

decisions—systematically vary with exogenous factors such as the order of survey responses.

Our paper contributes to the literature on technology and service provision. Autor, Levy, and Murnane (2003) show that computerization in the second half of the 20th century did little to reduce the cost of nonroutine cognitive tasks, which are often required for personalized services. The digitalization of the 21st century, however, has raised hopes of delivering low-cost customized services at scale (Varian, 2010; Goldfarb and Tucker, 2019). In finance, Philippon (2020) shows that the cost of financial services has remained largely unchanged for centuries but argues that technologies such as robo-advising may reduce it. Recent work has examined the role of robo-advisors in helping retail investors overcome behavioral biases (D'Acunto, Prabhala, and Rossi, 2019; D'Acunto and Rossi, 2023; Reher and Sokolinski, 2024). Our paper identifies a potential weak link in technology-driven personalization: when such services rely on survey inputs, biases in those surveys may become embedded in the service itself and shape customers' decisions.

More broadly, our paper contributes to the emerging literature on how online decision environments shape financial investments. For example, Liao et al. (2021) document how user-interface design affects financial decisions. In our setting, the risk assessment survey is part of the user interface—and thus the decision environment—through which survey-induced biases can become embedded in the investment process.

The response-order effect has been extensively studied in psychology (Krosnick and Presser, 2010). Our paper extends this line of research to economic surveys and studies its effects on high-stakes real-world financial decisions. More broadly, prior studies document order effects beyond response order in a variety of settings, including auction revenue (Hong et al., 2015), attention to academic papers (Feenberg et al., 2017), investor attention (e.g., Hartzmark, 2015; Frydman and Wang, 2020; Jiang et al., 2022), citations (Huang, 2015; Brogaard et al., 2023), school admissions (Jurajda and Munich, 2010), and election outcomes (Chen et al., 2014).

The rest of the paper is as follows. Section 2 reports a study based on a survey of U.S. investors. Section 3 presents an RCT at a leading robo-advisor in China. Section 4 concludes.

2. A Nationwide Survey

We administered a nationwide survey through the American Association of Individual Investors (AAII), a nonprofit organization with around 160,000 members. According to the AAII website, a typical member is a male in his mid-60s with a bachelor's or graduate degree, and members tend

to be affluent, with a median portfolio size exceeding $1 million. AAII surveys are widely used in prior research: Greenwood and Shleifer (2014) show that expectations elicited from AAII surveys are highly correlated with those from other investor surveys, and Jiang, Peng, and Yan (2024) use AAII data to examine the relationship between personality traits and investment decisions.

We administered two versions of the survey. Both versions contained the same 17 questions in the same order; the only difference was the order of the response options. Each respondent was randomly assigned to one of the two versions with equal probability. The version shown in Appendix A1 is the "original-order survey," while in the "reverse-order survey" the response options are reversed for each question. For example, the first question is "How old are you?" The response options for the two versions are as follows:

Original order	*Reverse order*
18 – 39 years	70 years or older
40 – 49 years	60 – 69 years
50 – 59 years	50 – 59 years
60 – 69 years	40 – 49 years
70 years or older	18 – 39 years

Our survey was distributed via an email by the AAII to its members on October 27, 2020, and data collection concluded on November 2, 2020. We received responses from 3,146 individuals: 1,561 from the original-order survey and 1,585 from the reverse-order one.

Respondent *i*'s answer to survey question *j* is coded as

$$Choice_{ij} = k, \tag{1}$$

if the respondent selected the *k*-th answer in the original-order list. For example, if respondent *i* receives the original-order survey and selects the fourth response to the first question (i.e., "60–69 years"), the answer is coded as $Choice_{i1} = 4$. If respondent *i* receives the reverse-order survey and selects the second response (which is also "60–69 years"), the answer is likewise coded as $Choice_{i1} = 4$. That is, the answer "60–69 years" is coded as 4 in both survey versions.

Table 1 reports the summary statistics. The first three questions gather basic information. As expected, our sample is skewed towards seniors and males. Around 90% of the respondents are over 60 years old and 93% are male. The respondents are experienced investors; around 92% had more than 20 years of experience.

Our main focus is Questions 4 and 5, which are about investment horizon and risk tolerance and closely resemble the survey questions commonly used by robo-advisors. Around 60% of our

respondents select "More than 20 years" or "Beyond my lifespan (estate)." Around 63% indicate that they can tolerate a 30% or 40% loss within a year.

Questions 6 and 7 are about forecasting the GDP growth and S&P 500 stock index return, respectively, for the next 12 months. Questions 8 and 9 are about the respondent's experiences during the COVID pandemic. Questions 10 through 12 focus on their opinions on the COVID pandemic. Questions 13 through 15 are about their political leaning and the U.S presidential election. Finally, the responses to Question 16 show that 91% of our respondents have a bachelor's degree or above. The responses to Question 17 indicate that 69% live in suburbs and that rural and urban areas account for around 16 percent each. More details of this survey can be found in Xiao and Yan (2023), which examines how information from social network disproportionately shapes expectations across domains.

2.1 The Response-order Effect for Robo-advisory Questions

A common response-order effect is that respondents are more likely to select options near the top of the list (Krosnick and Presser, 2010). This pattern is consistent with the anchoring-and-adjustment heuristic (Tversky and Kahneman, 1974): individuals anchor on the first options they encounter and adjust insufficiently, biasing choices toward earlier options. This implies that, given the definition in (1), responses to the reverse-order survey should have higher values of $Choice$ than those to the original-order survey.

As shown in Table 1, the number of response options ranges from two to six. To ensure comparability across questions, we standardize $Choice_{ij}$ to a score between 0 and 100:

$$Score_{ij} = \left(Choice_{ij} - 1\right)\frac{100}{\#Choice_j - 1}, \tag{2}$$

where $\#Choice_j$ is the number of response options for Question j. For example, if a question has five response options, the corresponding scores are 0, 25, 50, 75, and 100 for choices 1 through 5, respectively. To estimate the response-order bias for question j, we run the following cross-sectional regression:

$$Score_{ij} = \alpha + \beta\, Reverse_i + \varepsilon_{ij}, \tag{3}$$

where $Reverse_i$ is a dummy variable equal to one if respondent i is assigned to the reverse-order survey and zero otherwise. The coefficient β captures the average difference in standardized scores between respondents assigned to the reverse-order and original-order surveys. Results are

reported in Panel A of Table 2.

The first two rows indicate that response-order effects are both statistically and economically significant for survey questions commonly used by robo-advisors. The first row corresponds to Question 4—"For the goal(s) of your overall investment portfolio, what is your approximate time horizon?"—for which the estimated coefficient on Reverse is 10.67 (s.e. = 0.93). The second row corresponds to question 5—"Recognizing that investments can fluctuate in value, for a $100,000 investment, how much of a decline in the value of your investment portfolio could you tolerate over the course of a year?"—with an estimated coefficient of 9.08 (s.e. = 1.14). Given the 0–100 scaling of the score, these estimates correspond to shifts of roughly 9–11 percentage points, indicating economically large changes in responses.

We can also assess the magnitude of the response-order effect by examining respondents' choices. For Question 4, 50% of respondents select the longest investment horizon ("Beyond my lifespan (estate)") when it appears at the top on the reverse-order survey, compared with 35% when it appears at the bottom in the original-order survey ($p < 10^{-16}$). For Question 5, the most risk-tolerant option is chosen by 39% of respondents in the reverse-order survey, compared with 28% in the original-order survey ($p = 6.29 \times 10^{-10}$). These differences—on the order of 10 to 15 percentage points—underscore the economic importance of the response-order effect.

2.2 Objective vs Subjective Questions

We conjecture that the response-order effect is stronger when respondents are more uncertain about their answers. This is because anchoring effects are stronger when respondents are less confident (Jacowitz and Kahneman, 1995). Respondents are generally more certain when answering objective questions—such as demographic characteristics—than subjective ones, such as risk tolerance. In our survey, objective questions include those on age, gender, and investment experience (Questions 1–3), recent personal experiences during the pandemic (Questions 8–9), and education and residence (Questions 16–17), for which we expect little or no order effect. The remaining questions elicit subjective judgments, for which respondents may be uncertain which option to choose. We therefore expect stronger response-order effects for these questions.

Consistent with this hypothesis, Panel A of Table 2 shows strong evidence of the response-order effect for subjective questions. For nine of the ten questions, the coefficient on $Reverse$ is positive and highly statistically significant. In contrast, for six of the seven objective questions,

the estimated coefficient on $Reverse$ is statistically indistinguishable from zero.[2]

We also compare the effects for subjective and objective questions using panel regressions. Specifically, we pool all survey questions and estimate regression (3), with results reported in the first two columns of Panel B of Table 2. Column (1) shows that the coefficient on $Reverse$ is 2.54 ($s.e.$ = 0.29), implying that, on average across questions, reversing the response order increases Score by 2.54 points. In column (2), we include question fixed effects; the coefficient on Reverse remains similar in both magnitude and statistical significance. To contrast subjective and objective questions, we augment regression (3) with an interaction term $Reverse_i \times Sub_j$, where Sub_j is a dummy variable equal to one if question j is subjective and zero otherwise. The results, reported in columns (3) and (4), show that the response-order effect is concentrated in subjective questions. In column (4), the coefficient on $Reverse$ is 0.36 ($s.e.$ = 0.47), indicating no statistically significant effect for objective questions. In contrast, the effect for subjective questions is substantially larger, with a point estimate of 4.11 (= 3.75 + 0.36).[3]

Note that there are no financial stakes involved, and respondents are not compensated for their participation. One potential concern is that respondents may exert limited effort, for example by selecting options at the top of the list, which could mechanically generate a response-order effect. However, the cross-sectional patterns in our data are not consistent with this low-effort interpretation.

A more important concern is whether these results generalize to settings with financial stakes. In a robo-advisory context, where survey responses have real consequences, investors may behave differently. Do response-order effects persist in such settings, and do they influence actual investment decisions? We address these questions next.

---

[2] The two exceptions are worth noting. Question 10—whether face masks can prevent COVID infection for the wearer—is a subjective question, but respondents appear quite certain about their answers: approximately 84% report being "probably" or "definitely" confident about their answers, suggesting limited scope for response-order effects. Question 3—on investment experience—is notable because, although objective, it requires recalling the distant past: 92% of respondents report more than 20 years of experience. They may not recall precisely, leaving scope for a response-order effect.

[3] In the Online Appendix, we quantify the response-order effect on expectations about the macroeconomy, public health, and politics. We find that its magnitude is comparable to that of two prominent events at the time of our survey. First, the survey was administered in the early stage of the COVID-19 pandemic, when there was substantial uncertainty about its trajectory. Second, it was conducted immediately before the 2020 presidential election, when political divisions were pronounced and uncertainty about the outcome was high. The response-order effect is therefore comparable to the impact of a global pandemic and political leanings during one of the most divisive elections in recent history.

3. A Pre-registered RCT

Motivated by the survey evidence above, we collaborate with a leading robo-advisor in China (hereafter, the "Advisor") to conduct a randomized controlled trial (RCT) examining the response-order effect in a setting with substantial financial stakes. The RCT was preregistered with the AEA RCT Registry (AEARCTR-0014422).

3.1 Background

Established in 2014, the Advisor is one of the first robo-advisors in China. As of 2024, it serves approximately 150,000 clients and manages RMB 20 billion in assets. Similar to robo-advisors in the U.S., the Advisor assesses clients' investment preferences through a survey. When new clients open accounts with the Advisor, they are required to complete a risk assessment survey (see Appendix A2) consisting of 11 questions covering financial literacy, investment experience, income, risk tolerance, liabilities, and related factors. Typically, existing clients are prompted to retake the survey every two years.

For each survey question, a client is assigned a numerical score based on his answer. The sum of these scores across all 11 questions constitutes the client's overall *risk score*, which ranges from 18 to 89. Clients are assigned to one of five *risk categories* based on their risk scores: C1 (18–19), C2 (20–28), C3 (29, 38), C4 (39–68), and C5 (69–89). There is one exception to this classification rule: Question 8 of the survey asks, "What is your investment preference?", and the listed responses in the original-order survey are as follows:

4% return with no principal loss
6% return with a maximum 2% potential principal loss
10% return with a maximum 7% potential principal loss
15% return with a maximum 15% potential principal loss
30% return with a maximum 40% potential principal loss

If a client selects the first choice—"4% return with no principal loss"—his risk score will be capped at 19, and he will be assigned to category C1.

The Advisor offers a range of financial products, and our analysis focuses on two main types. We refer to the first type as "self-directed" investment products to highlight users' active role in the investment process. This category includes mutual funds, money market funds, and wealth management products offered by third parties, with the Advisor serving as the intermediary.

The Advisor also offers portfolios composed of these third-party products, allowing clients to invest directly in those portfolios rather than individual components.

The Advisor classifies these products into five risk categories, R1 through R5, with R5 being the highest risk level. For compliance reasons, a user is only allowed to purchase products whose risk levels are at or below his assigned risk category. For example, a client with the risk category C4 can invest in products rated R1 through R4. If the client attempts to purchase an R5 product, he would be reminded that the product's risk level is beyond his category and prompted to retake the risk survey before proceeding.[4]

The main page of the Advisor's app displays a curated set of products based on a proprietary algorithm that incorporates a client's risk score and investment history. According to the management at the Advisor, while these prominently displayed products tend to receive more attention, users can easily navigate away to explore and purchase other offerings.

The second type is "robo-advised" investment products, which are similar to the robo-advisory services in the U.S. market. The Advisor provides three robo-advised investment strategies: "Aggressive," "Balanced," and "Stable" strategies, which have high, medium, and low risk levels, respectively. Each strategy aims to build and maintain a target portfolio tailored to a specific risk level. For example, as shown in Panel A of Table 3, the target allocation to Chinese and U.S. equities is 47% under the Aggressive strategy but only 24% under the Stable strategy.

To invest in a strategy, a user must set up a dedicated account. A user may open accounts for one or multiple strategies, but the risk assessment may steer him toward a particular one. For example, a C5 user's report begins with "Your risk profile is Aggressive," followed by a description of how Aggressive investors typically invest. Notably, the label "Aggressive" is used for both the risk profile and the corresponding strategy, potentially steering the user toward the Aggressive strategy.

If a C4 user attempts to set up an account for the Aggressive strategy, he is warned that the product's risk level exceeds his assigned category. Although the user can ignore the warning and proceed, compliance restrictions may still limit subsequent investment choices. In particular, he can follow the strategy's recommendations only if they do not involve purchasing R5 products. If

[4] In principle, a user can contact customer service to declare that he understands the risks associated with an R5 product and insist on purchasing it without retaking the risk-assessment survey. In practice, however, this is extremely rare. In our sample, only 19 users purchased R5 products while classified in category C4.

a recommended transaction includes the purchase of an R5 product, he is again warned that the product exceeds his risk category and must retake the survey to obtain a C5 classification before proceeding.

We note that users play a more active role in self-directed investments than in robo-advised investments. A user needs to actively make all decisions in his self-directed investments. In contrast, for robo-advised investments, the trading strategies are designed by the Advisor. Although a user needs to authorize each recommended transaction, he can only accept or reject a recommended transaction, but cannot make an alternative transaction. It is also noteworthy that users in the U.S. robo-advisory market play an even more passive role. Typically, once a user completes the risk-assessment survey and establishes an account with a robo-advisor, subsequent investment decisions are executed automatically based on the user's risk-assessment results.

3.2 The RCT Procedure

As in the AAII survey, treatment is implemented by randomizing the response order for each survey question. The risk assessment survey was presented in two versions that differ only in the order of listed responses. In one version—referred to as the original-order survey and shown in Appendix A2—the responses for each question are ordered from lowest to highest risk score. In the other version—referred to as the reverse-order survey—the responses are listed from highest to lowest risk score. Each user was randomly assigned to one of the two versions with equal probability.

Our RCT began on October 25, 2024. We refer to users with accounts prior to this date as "existing users," and those who opened accounts afterward as "new users." For new users, treatment was administered from October 25 to December 27, 2024. During this period, when an individual opened an account, he first obtained an account number and then completed the risk assessment survey. Survey assignment was determined by account-number parity: users with even numbers received the original-order survey, while those with odd numbers received the reverse-order survey.

For existing users, the treatment was rolled out in two waves. On November 1, 2024, 5,000 users with holdings in robo-advised products were randomly selected and invited to retake the risk assessment survey upon logging in. If a user chose to proceed, survey assignment followed the same odd–even account-number rule as for new users. Users could ignore the invitation, but the

prompt reappeared upon subsequent logins and disappeared after three declines. A second wave of 5,000 randomly selected users was invited on December 18, 2024, under the same design, with no overlap between the two groups.

The trial concluded on April 1, 2025, when the data were retrieved. As shown in Panel B of Table 3, our *RCT sample* comprises 7,141 users. The vast majority fall into categories C4 and C5, with only 260 users classified into categories C1 through C3. The second and third rows show that 3,542 users received the original-order survey, while 3,599 received the reverse-order version.

The RCT for new users ended on December 27, 2024. When retrieving the data on April 1, 2025, we also obtained records for 3,142 new users who opened accounts between December 28, 2024, and March 31, 2025, after the odd–even rule no longer determined survey assignment. All of these users received the original-order survey, and we refer to them as the *placebo sample*.

Panel C reports summary statistics. For the RCT sample, the average risk score is 67, the average age is 46, and 60% of users are male. During the sample period, 1,746 users made robo-advised investments, with an average purchase of RMB 27,830, while 1,419 users made self-directed investments, with an average purchase of RMB 54,000. The placebo sample consists of new users who opened accounts between December 28, 2024, and March 31, 2025. These users have lower average risk scores (61) and are younger, with an average age of 40.

Panel D reports summary statistics separately for existing and new users in the RCT sample. The sample comprises 5,818 existing users and 1,323 new users. Prior to the RCT, existing users held an average of RMB 315,830 in total assets.

Panel E shows that the original- and reverse-order groups are well balanced across all observable characteristics, including gender, age, new-user status, and pre-treatment risk tolerance and asset holdings. None of these differences are statistically significant.

3.3 Order Effects on Risk Assessments

Does the order effect persist when large financial stakes are involved? To examine this, we estimate the following regression using our RCT sample:

$$Y_i = \alpha + \beta Reverse_i + \gamma X_i + \varepsilon_i, \quad (4)$$

where Y_i is user i's risk score or category, $Reverse_i$ is a dummy variable equal to one if user i's account number is odd (and hence the user received the reverse-order survey), and zero otherwise. X includes control variables: the user's age, gender, and initial asset holding at the Advisor.

The regression results, reported in Panel A of Table 4, indicate a strong order effect. In column (1), the coefficient on $Reverse$ is 3.35 (*s.e.* = 0.31), suggesting that reversing the response order increases the average risk score by 3.35. Column (2) adds control variables, and the coefficient on $Reverse$ remains mostly unchanged.

Does response order influence users' risk categories? Panel B of Table 3 provides a clear answer. In category C1, the number of users drops from 161 in the original-order group to 78 in the reverse-order group ($p = 2.29 \times 10^{-8}$). At the upper end, response order also shifts the share of users in category C5. Specifically, the share rises from 47.9% (1,697 users) in the original-order group to 56.5% (2,035 users) in the reverse-order group—an increase of 8.6 percentage points ($p = 2.83 \times 10^{-13}$). Thus, reversing the response order shifts users away from the lowest risk category and toward the highest, with economically large effects at both tails.

We also estimate regression (4) by setting Y_i to 1 through 5 if respondent i's risk category is C1 through C5, respectively. Column (3) of Panel A shows that the coefficient on $Reverse$ is 0.16 (s.e. = 0.02), indicating that—consistent with the results based on risk scores—reversing the response order increases a user's risk category. Column (4) adds control variables, and the coefficient on $Reverse$ remains mostly unchanged.

Finally, we conduct a falsification test using the placebo sample. Recall that response order is no longer randomized in this sample, as these users opened their accounts after the RCT concluded. Specifically, we estimate regression (4) on this sample, with $Reverse_i$ defined as before: it equals to one if user i's account number is odd and zero otherwise. In the absence of randomization, we expect no systematic difference in risk assessments between users with $Reverse_i = 1$ and those with $Reverse_i = 0$. Indeed, as shown in Panel B, the estimated coefficients on *Reverse* are statistically insignificant across all specifications. It is worth noting that this null result is not due to a lack of statistical power: the standard errors are small, and the estimated coefficients are close to zero.

Together, these results support the identification strategy by showing that the odd–even number rule has no predictive power in the absence of randomization. In contrast, when assignment is randomized, the same rule generates economically and statistically significant differences in risk assessments. That is, despite the large financial stakes involved, the response order still has a substantial impact on users' reported risk preferences.

3.4 Why Should the Response Order Affect Investment Decisions?

The previous subsection shows that the response order influences users' risk scores and risk categories. But does it also affect their investment decisions? There are at least two reasons to conjecture that it does. First, the order effect on risk assessments may influence users' investment decisions through a *nudging channel*. Since the risk assessments are used as an input for a robo-advisor's algorithm for its investment recommendations, the order effect becomes embedded in those recommendations, potentially nudging users to act. In the U.S., for example, once users complete their risk assessment surveys and set up accounts with a robo-advisor, the rest of the investment decisions are executed automatically based on the risk assessment results (Reher and Sokolinski, 2024). Hence, the order effect on risk assessments automatically turns into an effect on users' investment decisions. As noted in Section 3.1, investors play a more active role in the robo-advisory market in China. Hence, one would expect the order effect on investment decisions to be weaker in our setting than in the U.S. robo-advisory market.

Second, the order effect on risk assessments may also influence users' investment decisions through a *compliance channel*. For regulatory reasons, a user at the Advisor can only purchase assets with risk ratings at or below his assigned risk category. For instance, a user classified as C4 may purchase products rated R1 through R4. If the user attempts to purchase an R5 product, he would be reminded that the product's risk rating exceeds his assigned risk category. To proceed with the purchase, the user needs to retake his risk assessment survey to obtain a C5 assessment. As a result of this restriction, we expect the response order to influence users' investment choices through its effects on their risk categories.

Let us take C1 investors as an example. Due to their low stated risk tolerance, the Advisor prohibits these investors from purchasing any products with exposure to the equity market. As reported in Table 3, 4.5% (161 out of 3,542) of users in the original-order group are classified as C1, compared with 2.2% (78 out of 3,599) in the reverse-order group ($p = 2.29 \times 10^{-8}$). Thus, exposure to the original-order survey increases the share of C1 users by 2.3 percentage points, implying that an additional 2.3% of users are effectively excluded from equity market participation. To the extent that some of these users would choose to invest in equities if assigned a higher risk category, response order has real effects on investment decisions.

A symmetric pattern emerges at the upper end of the risk distribution. As shown in Table 3, the share of users in the highest risk category, C5, decreases from 56.5% (2,035 out of 3,599

users) in the reverse-order group to 47.9% (1,697 out of 3,542 users) in the original-order group ($p = 2.83 \times 10^{-13}$), a decline of 8.6 percentage points. This shift implies that a substantial fraction of original-order users are moved away from C5 category. To the extent that some of these users would otherwise prefer R5 products, response order again has real effects on investment decisions by steering them toward more conservative allocations.

In the following, we directly analyze the response-order effects on users investment decisions during the trial. As noted in Section 3.1, self-directed investments require more active involvement from users, whereas robo-advised investments resemble the U.S. robo-advisory market, where investors play a relatively passive role. Accordingly, we expect the order effect on investment decisions to be stronger for robo-advised than for self-directed investments.

3.5 Order Effects on Robo-advised Investments

This section examines the order effect on robo-advised investments. We first analyze the effects along the intensive margin. That is, we examine how the response order influences users' investment allocations, conditional on their participation. As shown in Table 4, the reverse-order survey increases users' assigned risk categories. How might this shift affect their investment decisions? Intuitively, the order effect can operate through both a nudging mechanism and a compliance mechanism.

First (nudging channel), the reverse order users' higher risk categories may guide them towards riskier products. This is especially the case for robo-advised investment products. For example, a user classified to the C5 category in the risk assessment survey receives the following report:

> *Your risk profile is "Aggressive." Aggressive investors typically seek to achieve above-average returns and have a strong tolerance for market volatility and potential capital losses. Their investment objective is to earn high returns while accepting high levels of risk.*
>
> *Most of the assets for aggressive investors are allocated to domestic and international equities and other similar products. A very small portion is allocated to cash and fixed-income products, adjusted based on market conditions, to help diversify risk.*

Notably, the term "Aggressive" is also used to label the investment strategy, reinforcing the alignment between the assigned risk profile and the recommended strategy.

The operational procedure may also nudge users toward investment strategies aligned with their assigned risk categories. Specifically, when a user opens a special purpose account for robo-advised investments, the assigned risk category can influence the selection of investment strategies. For instance, the Aggressive strategy is designed for C5 users. If a C4 user attempts to enroll in this strategy, a warning message pops up and indicates that the strategy's risk level exceeds the user's risk category. Although the user can override the warning and proceed, such prompts are likely to deter some C4 users from pursuing the Aggressive strategy.

Second (compliance channel), the compliance mechanism may further steer reverse-order users—who are more likely to receive higher risk categories—toward riskier strategies. As noted earlier, a C4 user is permitted to invest in the Aggressive strategy. However, compliance constraints may limit the size of that investment. For example, if the Aggressive strategy recommends purchasing an R5 product, the user will be prevented from completing the transaction unless he retakes the risk assessment survey and obtains a C5 risk category. This restriction effectively reduces a C4 user's investment size in the Aggressive strategy.

To estimate the order effect on investments in robo-advised products along the intensive margin, we focus on investors who made purchases during the sample period. We estimate regression (4) with the dependent variable A_i, defined as user i's total purchases of the Aggressive strategy during the sample period. The results, reported in column (1) of Panel A of Table 5, show that the coefficient on $Reverse$ is 3.28 (s.e. = 1.66). This implies that exposure to the reverse-order survey increases investment in the Aggressive strategy by approximately RMB 3,280 on average—about 12% of the average purchase size (RMB 27,830), indicating a sizable effect along the intensive margin.

Our hypothesis is that response-order effects operate through users' risk categories: response order influences the assigned category, which in turn steers investment choices. To test this mechanism, we decompose a user's risk category into two components and examine their effects on investment. The first component is driven by an exogenous factor—the response order—while the second captures the user's underlying risk tolerance.

Specifically, since all users in this sample fall into categories C4 and C5, we estimate regression (4) with $C5_i$ as the dependent variable, where $C5_i$ is an indicator equal to one if user i is classified into the C5 category. Column (2) shows that the coefficient on $Reverse$ is 0.07 (*s.e.* = 0.02), implying that exposure to the reverse-order survey increases the probability of being

classified into the C5 category by 7.40 percentage points. Based on this regression, we decompose $C5_i$ into its fitted value $\widehat{C5}_i$ and residual $C5Residual_i$. The fitted value captures variation driven by the exogenous response-order manipulation, while the residual reflects the user's underlying risk tolerance.

Column (3) includes both components in a regression of A_i, user i's total purchases of the Aggressive strategy during the trial. The coefficient on $\widehat{C5}_i$ is 44.26 (*s.e.* = 22.28), implying that a one-percentage-point increase in the probability of being classified into the C5 category due to the exogenous manipulation increases the user's purchase of the Aggressive product by RMB 443 during the trial. In contrast, the coefficient on $C5Residual_i$ is 8.90 (s.e. = 1.35), indicating that, as expected, users with higher underlying risk tolerance invest more in the Aggressive strategy.

We now turn to the extensive margin, examining whether response order affects users' participation in robo-advised investments. Specifically, we estimate regression (4) on the full RCT sample, using $DRobo_i$ as the dependent variable, which equals one if user i purchased robo-advised products during the trial and zero otherwise. The results are reported in Panel B of Table 5. Column (1) shows that the coefficient on $Reverse$ is –0.0254 (*s.e.* = 0.0105), indicating that users assigned to the reverse-order survey are 2.54 percentage points less likely to invest in robo-advised products during the trial.

This pattern cannot be explained by the compliance channel. If compliance were the primary driver, reverse-order users—who tend to receive higher risk categories—should face fewer investment restrictions and therefore be more likely to invest.

One possible explanation is that nudging backfired in this case. By inflating risk categories, the order effect steers some users toward riskier strategies that they may perceive as too risky, leading them to opt out of investing altogether. Columns (2) and (3) provide evidence consistent with this interpretation. As in Panel A, column (2) decomposes $C5_i$ into a component driven by the exogenous response-order manipulation, $\widehat{C5}_i$, and a residual component, $C5Residual_i$, that reflects the user's underlying risk tolerance. Column (3) shows that the coefficient on $C5Residual_i$ is 0.10 (*s.e.* = 0.01), indicating that users with higher underlying risk tolerance are more likely to invest. In contrast, the coefficient on $\widehat{C5}_i$ is –0.30 (s.e. = 0.12), suggesting that an exogenously inflated risk category reduces the likelihood of investment.

Together, these results suggest that the response-order effect can reduce participation by pushing some users into risk categories that do not match their true risk tolerance. However,

conditional on investing, an inflated risk category induced by response order leads users to allocate more to the Aggressive strategy.

3.6 Order Effects on Self-directed Investments

We next examine whether response order affects self-directed investments by applying the same analysis as in Table 5 to the sample of self-directed trades. To estimate the effect along the intensive margin, we estimate regression (4) using the sample of users who purchased self-directed investment products during the trial, and the results are reported in Panel A of Table 6. In column (1), the dependent variable is a user's total purchase of R5 self-directed products during the trial. The coefficient on $Reverse$ is 3.09 (*s.e.* = 1.84), implying that exposure to the reverse-order survey increases investment in R5 products by approximately RMB 3,090 on average, which corresponds to about 6% of the average purchase size of self-directed products (RMB 54,000). This effect is similar to, but somewhat weaker than, that for robo-advised products, both statistically and economically.

Columns (2) and (3) conduct a similar decomposition of users' risk assessments. Column (2) decomposes $C5_i$ into an exogenous component driven by the response-order effect, $\widehat{C5}_i$, and a residual component, $C5Residual_i$, that reflects the user's underlying risk tolerance. Column (3) shows that a user's total purchase of R5 products is increasing with both components.

To estimate the effect along the extensive margin, we estimate regression (6) on the full RCT sample, using $DSelf_i$ as the dependent variable, which equals one if user i purchased any self-directed products during the trial and zero otherwise. The results are reported in column (1) of Panel B of Table 6. The coefficient on Reverse is –0.0178 (s.e. = 0.0090), indicating that users assigned to the reverse-order survey are 1.78 percentage points less likely to invest in self-directed products. Again, this effect is weaker, both statistically and economically, than its counterpart for robo-advised investments. We conduct a similar decomposition analysis in column (2), and column (3) show that users with higher underlying risk tolerance are more likely to invest in self-directed products, while an exogenously inflated risk category—driven by the response-order effect—reduces the likelihood of investment.

In summary, the analysis in this section suggests that reversing the response order reduces users' likelihood of investing in self-directed products. However, conditional on investing, reverse-order users allocate more to products in the riskiest category, R5. The effects along both

the intensive and extensive margins are weaker than their counterparts for robo-advised investments. This pattern is consistent with our earlier observation that response-order effects are likely attenuated in self-directed settings, where users play a more active role in their investment decisions.

3.7 Implications

Our findings raise an important ethical concern. When a robo-advisor's incentives are not fully aligned with those of its clients, it may benefit from steering users toward products that are not optimal from the clients' perspective. Because investor preferences are inferred from survey responses that can be influenced by design choices such as response order, a robo-advisor could exploit this mechanism to guide users toward products that serve its own interests—while maintaining the appearance that users "prefer" those products.

More broadly, our findings have implications for constructs such as risk aversion and preferences. These concepts are deeply embedded in economic analysis and are often treated as primitives. Our evidence suggests that they are not fully determinant in the sense that they cannot be inferred from either surveys or revealed preferences. This is because response order affects not only the stated preferences in the survey but also actual investment decisions. As a result, an external factor—response order—can shape the inferred preference parameters even under a revealed-preference approach, challenging the view that such parameters are intrinsic and context-independent.

How can response-order bias be mitigated? The appropriate solution depends on the objective. If the goal is to estimate average opinions in a population, the bias can be mitigated by randomizing response order—for example, assigning half of respondents to the original order and half to the reverse order—so that the biases largely offset in the aggregate. If the goal is to measure an individual's preferences, a different approach is needed. For instance, a robo-advisor seeking to estimate a client's risk tolerance could include two differently phrased questions on risk tolerance with opposite response orders. Averaging the inferred risk tolerance across these two questions would yield a more accurate estimate, as the two biases offset each other.

4. Conclusions

Our evidence shows that FinTech can transmit a seemingly innocuous psychological bias into the

two-trillion-dollar robo-advisory market. We begin with a nationwide survey of more than 3,000 experienced U.S. investors and document a substantial response-order effect in the risk-tolerance questions commonly used by robo-advisors. Building on these findings, we partner with a leading robo-advisor in China to conduct a preregistered RCT involving more than 10,000 investors. We find that, even with substantial financial stakes, the order of survey responses continues to exert a significant influence on clients' risk profiles, which directly shape product recommendations. More importantly, response order also has economically meaningful effects on users' actual investment behavior.

While our evidence is based on an RCT in China, the response-order effect may be even stronger in the U.S., where investors often play a more passive role. Once users establish and fund their accounts with a robo-advisor, subsequent investment decisions are typically executed automatically based on survey inputs. As a result, any bias introduced by response order can be transmitted directly into investment decisions.

Our findings raise an ethical concern: a robo-advisor may steer users toward products that advance its own interests while preserving the appearance that these choices reflect users' preferences. More broadly, the evidence suggests that preference constructs such as risk aversion are not fully determinate: their estimated values—whether inferred from surveys or from observed investment decisions—systematically vary with exogenous factors such as response order.

Finally, our analysis focuses on short-term effects. A longer-term RCT would be needed to assess the persistence of these effects on investment behavior and to provide sufficient statistical power to evaluate investment performance. We leave these important questions for future research.

References

Autor, David, Frank Levy, and Richard Murnane, 2003, The Skill Content of Recent Technological Change: An Empirical Exploration, *Quarterly Journal of Economics* 118, 1279–1333.

Baumol, William, 1967, Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis." *American Economic Review* 57, 415–426.

Brogaard, Jonathan, Joseph Engelberg, Sapnoti Eswar, and Edward Van Wesep, 2023, On the Causal Effect of Fame on Citations, *Management Science* 70, 7187–7214.

Chen, Eric, Gabor Simonovits, Jon A. Krosnick, and Josh Pasek, 2014, The impact of candidate name order on election outcomes in North Dakota, *Electoral Studies* 35, 115–122.

D'Acunto, Francesco and Alberto Rossi, 2023, Robo-Advice: Transforming Households into Rational Economic Agents, *Annual Review of Financial Economics* 15, 543-563.

D'Acunto, Francesco, Nagpurnanand Prabhala, Alberto Rossi, 2019, The promises and pitfalls of robo-advising, *Review of Financial Studies* 32, 1983–2020.

Feenberg, Daniel R., Ina Ganguli, Patrick Gaule and Jonathan Gruber, 2017, It's Good to Be First: Order Bias in Reading and Citing NBER Working Papers, *Review of Economics and Statistics* 99, 32–39.

Frydman, Cary and Baolian Wang, 2020, The Impact of Salience on Investor Behavior: Evidence from a Natural Experiment, *Journal of Finance* 75, 229–276.

Goldfarb, Avi, and Catherine Tucker, 2019, Digital Economics, *Journal of Economic Literature* 57, 3–43.

Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of Returns and Expected Returns, *Review of Financial Studies* 27, 714–746.

Hartzmark, Samuel, 2015, The Worst, The Best, Ignoring All the Rest: The Rank Effect and Trading Behavior, *Review of Financial Studies* 28, 1024–1059.

Hong, Harrison, Ilan Kremer, Jeffrey Kubik, Jianping Mei and Michael Moses, 2015, Ordering, Revenue and Anchoring in Art Auctions, *RAND Journal of Economics* 46, 186–216.

Huang, Wei, 2015, Do ABCs Get More Citations than XYZs? *Economic Inquiry* 53, 773–789.

Jacowitz, Karen E. and Daniel Kahneman, 1995, Measures of Anchoring in Estimation Tasks, *Personality and Social Psychology Bulletin* 21, 1161–1166.

Jiang, Lei, Jinyu Liu, Lin Peng, and Baolian Wang, 2022, Investor Attention and Asset Pricing Anomalies, *Review of Finance* 26, 563–593.

Jiang, Zhengyang, Cameron Peng and Hongjun Yan, 2024, Personality Differences and Investment Decision-Making, *Journal of Financial Economics* 153, 103776.

Jurajda, Stepan, and Daniel Munich, 2010, Admission to Selective Schools, Alphabetically, *Economics of Education Review* 29, 1100–1109.

Krosnick, Jon and Presser, Stanley, 2010, Question and Questionnaire Design, in *Handbook of Survey Research* (2nd Edition) James D. Wright and Peter V. Marsden (Eds). San Diego, CA: Elsevier.

Liao, Li, Jia Xiang, Zhengwei Wang, Hongjun Yan, and Jun Yang, 2021, User Interface and First-hand Experience in Retail Investing, *Review of Financial Studies* 34, 4486–4523.

Philippon, Thomas, 2020, On fintech and financial inclusion, working paper.

Reher, Michael and Stanislav Sokolinski, 2024, Robo advisors and access to wealth management, *Journal of Financial Economics* 155, 103829

Tversky, Amos and Daniel Kahneman, 1974, Judgment under Uncertainty: Heuristics and Biases, *Science* 185, 1124–1131.

Varian, Hal, 2010, Computer Mediated Transactions, *American Economic Review: Papers and Proceedings* 100, 1–10.

Xiao, Yaqing and Hongjun Yan, 2023, Relatable Information Disproportionately Shapes Expectations, Working paper.

Table 1. Summary Statistics for the AAII Survey

This table reports the summary statistics for each survey question. The survey was sent out by the American Association of Individual Investors (AAII) to its members on October 27, 2020; data collection stopped on November 2, 2020. The survey received responses from 3,146 individuals. The survey is presented in two different versions. Each respondent is randomly assigned one version with a 50% chance. Both versions have the same set of questions presented in the same order. The only difference is the order of response choices. 1,561 respondents faced the original-order survey (shown in the Appendix A1) and 1,585 faced the survey with the choices for each question in the reverse order. The first column lists the survey questions. Those with asterisks are classified as "subjective questions" and all others are "objective questions." The second column reports the number of choices (#choice) for each question. The next three columns report the mean, standard deviation of *Answer* (defined in equation (1)), and number of respondents (N) for each question for the original-order subsample. Similarly, the last three columns report the results for the reverse-order subsample.

		Original order (1,561 respondents)			Reverse order (1,585 respondents)		
Question	#choice	Mean	S.D.	N	Mean	S.D.	N
1 Age	5	4.43	0.77	1549	4.45	0.74	1567
2 Gender	2	1.93	0.26	1529	1.94	0.24	1544
3 Experience	4	3.86	0.49	1557	3.90	0.44	1576
4* Horizon	6	4.47	1.40	1553	5.01	1.21	1579
5* Risk tolerance	4	2.75	0.98	1550	3.02	0.93	1576
6* GDP	5	3.29	0.89	1548	3.44	0.78	1570
7* Stock return	5	3.32	0.71	1536	3.41	0.69	1564
8 COVID mild	2	1.50	0.50	1551	1.49	0.50	1577
9 COVID serious	2	1.27	0.45	1551	1.27	0.45	1575
10* Facemask	5	4.26	1.04	1555	4.24	1.08	1577
11* Vaccine forecast	5	2.91	0.79	1546	3.09	0.83	1572
12* Vaccine take	5	4.17	0.93	1551	4.24	0.92	1570
13* Political leaning	5	3.10	1.32	1545	3.23	1.34	1565
14* Vote	5	2.64	1.79	1509	2.78	1.83	1521
15* Election forecast	5	2.52	0.98	1533	2.64	0.99	1556
16 Education	6	5.44	0.84	1553	5.43	0.86	1576
17 Residence	3	1.99	0.55	1551	2.01	0.57	1576

Table 2. Response-order Effects

Panel A reports the results from the following regression:

$$SAnswer_{ij} = \alpha + \beta Reverse_i \times Sub_j + \delta Reverse_i + \gamma Sub_j + \varepsilon_{ij},$$

where $SAnswer_{ij}$ is respondent i's standardized answer to Question j and is defined in (3); $Reverse_i$ is 1 if respondent i faces a reverse-order survey and 0 otherwise; and Sub_j is 1 if Question j is a subjective question and 0 otherwise. Standard errors, in parentheses, are clustered by respondent. Then, for each question j, we regress $SAnswer_{ij}$ on $Reverse_i$, and Panel B reports the estimated coefficient on $Reverse$, its standard error, and the number of observations for each regression. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

Panel A. Response-order bias

		Reverse	s.e.	N
Subjective questions	4 Horizon	10.67***	0.93	3132
	5 Risk tolerance	9.08***	1.14	3126
	6 GDP	3.69***	0.75	3118
	7 Stock return	2.29***	0.63	3100
	10 Facemask	-0.50	0.95	3132
	11 Vaccine forecast	4.31***	0.73	3118
	12 Vaccine take	1.65**	0.83	3121
	13 Political leaning	3.34***	1.19	3110
	14 Vote	3.52**	1.65	3030
	15 Election forecast	2.94***	0.89	3089
Objective questions	1 Age	0.71	0.68	3116
	2 Gender	0.91	0.90	3073
	3 Experience	1.16**	0.56	3133
	8 COVID mild	-0.83	1.79	3128
	9 COVID severe	-0.04	1.60	3126
	16 Education	-0.21	0.61	3129
	17 Residence	0.80	1.00	3127

Panel B. Objective vs subjective questions

	(1)	(2)	(3)	(4)
Reverse	2.54***	2.56***	0.31	0.36
	(0.29)	(0.29)	(0.47)	(0.47)
Reverse × *Sub*			3.80***	3.75***
			(0.58)	(0.58)
Sub			-11.60***	-12.36***
			(0.42)	(0.70)
Constant	63.16***	48.69***	69.97***	49.80***
	(0.21)	(0.52)	(0.33)	(0.55)
Question FEs	No	Yes	No	Yes
N	52908	52908	52908	52908
R-squared	0.001	0.317	0.020	0.318

Table 3. Summary Statistics for the RCT

Panel A reports the composites of the target portfolios of the Aggressive-, Balanced-, and Stable-Strategies. Panel B reports the number of users in each risk category for our RCT Sample, its two subsamples (the Original-order Subsample and the Reverse-order Subsample), and the Placebo Sample. Panels C through E report the summary statistics for various subsamples. *Score* is a user's risk score as of April 1, 2025. *Category* is a user's risk category as of April 1, 2025, and is 1 through 5 if the user's is in category C1 through C5, respectively. *Reverse* is one if a user's account number is odd (and hence if the user is in the RCT, he receives the reverse-order survey), and zero otherwise. *Age* is a user's age, denominated in years. *Male* is a dummy variable that is one is a user is a male and zero otherwise. *BuyRobo* and *BuyDirect* denote a user's total purchases of robo-advised and self-directed investment products, respectively, during the RCT period. *InitialAll* is a user's total asset holding before entering the RCT. *InitialRobo* and *InitialDirect* denote a user's total holdings of robo-advised and self-directed investment products, respectively, prior to the RCT. *InitialScore* and *InitialCategory* are a user's risk score and category, respectively, before entering the RCT. Columns labeled by "N" report the number of observations. The last two columns report the difference and its standard error. *NewUser* is a dummy variable that is one is a user's account was opened before October 25, 2024, and zero otherwise.

Panel A. Target portfolios

	Aggressive	Balanced	Stable
Chinese Stocks (%)	35	25	20
U.S. Stocks (%)	12	8	4
Alternative Investments (%)	25	20	10
Money Market Funds (%)	18	12	6
Chinese Government Bonds (%)	10	35	60

Panel B. Sample Size

	C1	C2	C3	C4	C5	Total
RCT Sample	239	3	18	3149	3732	7141
Original-order	161	1	10	1673	1697	3542
Reverse-order	78	2	8	1476	2035	3599
Placebo Sample	256	0	28	1653	1205	3142

Panel C. RCT and Placebo Samples

	RCT Sample			Placebo Sample		
	Mean	S.D.	N	Mean	S.D.	*N*
Score	66.88	13.32	7141	61.25	16.59	3142
Category	4.42	0.81	7141	4.13	1.05	3142
Reverse	0.50	0.50	7141	0.49	0.50	3142
Age (years)	45.66	10.28	6440	39.69	11.12	1640
Male	0.60	0.49	6440	0.47	0.50	1640
BuyRobo (1000 RMB)	27.83	64.21	1746	19.56	18.33	21
BuyDirect (1000 RMB)	54.00	142.08	1419	23.39	46.87	1067

Panel D. Existing- and New-user Samples

	Existing Users			New Users		
	Mean	S.D.	*N*	Mean	S.D.	*N*
Score	67.95	12.15	5818	62.16	16.78	1323
Category	4.47	0.73	5818	4.18	1.05	1323
Reverse	0.50	0.50	5818	0.52	0.50	1323
Age (years)	46.47	9.99	5818	38.09	9.84	622
Male	0.61	0.49	5818	0.48	0.50	622
BuyRobo (1000 RMB)	27.84	64.40	1734	26.06	25.24	12
BuyDirect (1000 RMB)	51.41	146.26	959	59.41	132.93	460
InitialAll (1000 RMB)	315.83	658.54	5818			
InitialRobo (1000 RMB)	158.01	414.60	5818			
InitialDirect (1000 RMB)	87.65	293.16	5818			

Panel E. Original- and Reverse-Order Samples

	Original Order		Reverse Order			
	Mean	N	Mean	N	Difference	*s.e.*
*Male*100*	58.86	3189	60.47	3251	1.62	1.22
Age (year)	45.68	3189	45.65	3251	-0.04	0.26
NewUser	0.18	3542	0.19	3599	0.01	0.01
InitialScore	49.89	2903	49.47	2915	-0.43	0.54
InitialCategory	3.68	2903	3.65	2915	0.03	0.03
InitialAll (1000 RMB)	308.69	2903	322.94	2915	14.25	17.27
InitialRobo (1000 RMB)	155.16	2903	160.85	2915	5.69	10.87
InitialDirect (1000 RMB)	83.14	2903	92.15	2915	9.01	7.69
Score (post-treatment)	65.19	3542	68.53	3599	3.35	0.31
Category (post-treatment)	4.34	3542	4.50	3599	0.16	0.02
BuyRobo (1000 RMB)	25.70	910	30.14	836	4.44	3.08
BuyDirect (1000 RMB)	49.32	715	58.76	704	9.44	7.54

Table 4. Order Effects on Risk Assessment

This table reports the results from the following regression:

$$Y_i = \alpha + \beta Reverse_i + X_i + \varepsilon_i,$$

where $Reverse_i$ is a dummy variable equal to one if user *i*'s account number is odd (and hence if the user is in the RCT, he receives the reverse-order survey), and zero otherwise. Panels A and B are based on the RCT and Placebo Samples, respectively. For both panels, Y_i is user *i*'s risk score in columns (1) and (2); and is 1 through 5 if user *i*'s risk category is C1 through C5, respectively, for columns (3) and (4). X_i refers to all control variables, which are defined in Table 3. Standard errors are reported in parentheses. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

Panel A. RCT Sample

	Risk Score		Risk category	
	(1)	(2)	(3)	(4)
Reverse	3.35***	3.12***	0.16***	0.14***
	(0.31)	(0.28)	(0.02)	(0.02)
Male		5.00***		0.23***
		(0.30)		(0.02)
Age		-0.06***		-0.00***
		(0.02)		(0.00)
LnInitialAll		1.91***		0.10***
		(0.11)		(0.01)
NewUser		7.34***		0.48***
		(0.74)		(0.04)
Constant	65.19***	56.92***	4.34***	3.95***
	(0.24)	(0.88)	(0.01)	(0.05)
N	7,141	6,440	7,141	6,440
R-squared	0.02	0.11	0.01	0.08

Panel B. Placebo sample

	Risk Score		Risk Category	
	(1)	(2)	(3)	(4)
Reverse	0.41	-0.27	0.02	-0.03
	(0.59)	(0.62)	(0.04)	(0.04)
Male		3.52***		0.13***
		(0.63)		(0.04)
Age		0.17***		0.01***
		(0.03)		(0.00)
Constant	61.05***	56.32***	4.12***	4.09***
	(0.42)	(1.28)	(0.03)	(0.07)
N	3,142	1,640	3,142	1,640
R-squared	0.00	0.04	0.00	0.01

Table 5. Order Effects on Robo-Advised Investments

Panel A reports the analysis on the intensive margin. The sample includes users who made robo-advised investments during our sample period. Columns (1) and (2) report the results from the following regression:

$$Y_i = \alpha + \beta Reverse_i + X_i + \varepsilon_i,$$

where $Reverse_i$ is a dummy variable equal to one if user i's account number is odd (and hence the user receives the reverse-order survey), and zero otherwise. X_i refers to all control variables, which are defined in Table 3. The dependent variable Y_i is A_i, user i's total purchase of the *Aggressive* portfolio, denominated in 1000 RMB, in column (1), and is $C5_i$ in column (2), where $C5_i$ is a dummy variable that is one if user i's risk category is C5, and zero otherwise. The regression in column (3) the same as in column (1) with $Reverse_i$ replaced by $\widehat{C5}_i$ and $C5Residual_i$, which are the predicted value of $C5_i$ and residual based on the regression in column (2). Panel B reports similar analysis on the extensive margin and is based on the entire RCT Sample. In columns (1) and (3), the dependent variable Y_i is a dummy variable $DRobo_i$, which is one if user i invested in robo-advised investments during our sample period, and zero otherwise. Robust standard errors are reported in parentheses. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

Panel A. Intensive Margin

	A	*C5*	*A*
	(1)	(2)	(3)
Reverse	3.28**	0.0740***	
	(1.66)	(0.0221)	
Male	3.36*	0.18***	-4.82
	(1.80)	(0.02)	(4.41)
Age	0.11	-0.00	0.14
	(0.09)	(0.00)	(0.09)
InitialRobo	0.02***	0.00***	0.01**
	(0.00)	(0.00)	(0.00)
NewUser	8.92	-0.18	17.02**
	(8.22)	(0.15)	(8.30)
$\widehat{C5}$			44.26**
			(22.28)
C5Residual			8.90***
			(1.35)
Constant	-0.26	0.52***	-23.42*
	(4.68)	(0.06)	(12.39)
Observations	1,746	1,746	1,746
R-squared	0.05	0.05	0.06

Panel B. Extensive Margin

	DRobo	*C5*	*DRobo*
	(1)	(2)	(3)
Reverse	-0.0254**	0.0861***	
	(0.0105)	(0.0121)	
Male	0.01	0.17***	0.06**
	(0.01)	(0.01)	(0.02)
Age	-0.00***	-0.00*	-0.00***
	(0.00)	(0.00)	(0.00)
LnInitialAll	0.06***	0.03***	0.07***
	(0.00)	(0.00)	(0.00)
NewUser	-0.09***	0.06***	-0.08***
	(0.01)	(0.02)	(0.01)
$\widehat{C5}$			-0.30**
			(0.12)
C5Residual			0.10***
			(0.01)
Constant	0.19***	0.36***	0.29***
	(0.03)	(0.03)	(0.06)
Observations	6,440	6,440	6,440
R-squared	0.11	0.05	0.12

Table 6. Order Effects on Self-Directed Investments

Panel A reports the analysis on the intensive margin. The sample includes users who made self-directed investments during our sample period. Columns (1) and (2) report the results from the following regression:

$$Y_i = \alpha + \beta Reverse_i + X_i + \varepsilon_i,$$

where $Reverse_i$ is a dummy variable equal to one if user i's account number is odd (and hence the user receives the reverse-order survey), and zero otherwise. X_i refers to all control variables, which are defined in Table 3. The dependent variable Y_i is $R5_i$ (user i's total purchase of self-directed investment products in the risk category of R5, denominated in 1000 RMB) in column (1), and is $C5_i$ in column (2), where $C5_i$ is a dummy variable that is one if user i's risk category is C5, and zero otherwise, in column (2). The regression in column (3) the same as in column (1) with $Reverse_i$ replaced by $\widehat{C5}_i$ and $C5Residual_i$, which are the predicted value of $C5_i$ and residual based on the regression in column (2). Panel B reports similar analysis on the extensive margin and is based on the entire RCT Sample. In columns (1) and (3), the dependent variable Y_i is a dummy variable $DSelf_i$, which is one if user i invested in self-directed investments during our sample period, and zero otherwise. Robust standard errors are reported in parentheses. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

Panel A. Intensive Margin

	R5	*C5*	*R5*
	(1)	(2)	(3)
Reverse	3.09*	0.0751***	
	(1.84)	(0.0247)	
Male	2.04	0.15***	-4.09
	(1.58)	(0.03)	(4.32)
Age	0.15	-0.00	0.17
	(0.11)	(0.00)	(0.11)
InitialSelf	0.02	0.00***	0.02
	(0.02)	(0.00)	(0.02)
NewUser	15.35***	-0.11***	19.68***
	(3.70)	(0.03)	(5.78)
$\widehat{C5}$			41.22*
			(24.39)
C5Residual			9.41***
			(1.03)
Constant	-10.70	0.58***	-34.46*
	(7.40)	(0.06)	(20.00)
Observations	1,419	1,419	1,419
R-squared	0.08	0.06	0.09

Panel B. Extensive Margin

	DSelf	*C5*	*DSelf*
	(1)	(2)	(3)
Reverse	-0.0178**	0.0850***	
	(0.0090)	(0.0121)	
Male	-0.03***	0.18***	0.01
	(0.01)	(0.01)	(0.02)
Age	-0.00	-0.00	-0.00*
	(0.00)	(0.00)	(0.00)
LnInitialAll	0.05***	0.03***	0.06***
	(0.00)	(0.00)	(0.00)
NewUser	0.69***	0.04*	0.70***
	(0.02)	(0.02)	(0.02)
$\widehat{C5}$			-0.21**
			(0.11)
C5Residual			0.09***
			(0.01)
Constant	0.10***	0.36***	0.17***
	(0.02)	(0.03)	(0.05)
Observations	6,440	6,440	6,440
R-squared	0.24	0.05	0.25

Appendix A1. AAII Survey Questions

Question 1: How old are you?

18 – 39 years
40 – 49 years
50 – 59 years
60 – 69 years
70 years or older

Question 2: What is your gender?

Female
Male

Question 3: How long have you been investing?

0 – 9 years
10 – 15 years
16 – 20 years
More than 20 years

Question 4: For the goal(s) of your overall investment portfolio, what is your approximate time horizon?

Less than 1 year
1 - 5 years
6 - 10 years
11 - 20 years
More than 20 years
Beyond my lifespan (estate)

How certain do you feel about your answer to this question?

Very certain: I have a clear idea about my investment horizon
Somewhat certain: I have a rough idea about my investment horizon
Not certain: I have barely considered my investment horizon

Question 5: Recognizing that investments can fluctuate in value, for a $100,000 investment, how much of a decline in the value of your investment portfolio could you tolerate over the course of a year?

-10% (-$10,000)
-20% (-$20,000)
-30% (-$30,000)
-40% (-$40,000)

Question 6 What do you think GDP growth in the U.S. will be over the next 12 months?

Less than -5%
-5% to -2%
-2% to 2%
2% to 5%
Greater than 5%

How certain are you about your forecast?

I am very certain about my forecast
I am somewhat certain about my forecast
I am not certain about my forecast.

Question 7 What do you think the total return of the S&P 500 stock market index will be over the next 12 months?

Less than -15%
-15% to -5%
-5% to 5%
5% to 15%
Greater than 15%

How certain are you about your forecast?

I am very certain about my forecast
I am somewhat certain about my forecast
I am not certain about my forecast

Question 8 Do you directly know someone who has been infected by COVID-19 with mild or no symptoms?

No
Yes

Question 9 Do you directly know someone who has been severely ill due to COVID-19 infection?

No
Yes

Question 10 Do you believe that wearing a mask reduces *your* risk of COVID-19 infection?

Definitely not
Probably not
Not sure
Probably yes
Definitely yes

How certain are you about your answer?

I am very certain about my answer
I am somewhat certain about my answer
I am not certain about my answer

Question 11 When do you expect a COVID-19 vaccine to become widely available?

Within 3 months
3 months to 6 months
6 months to 12 months
1 year to 3 years
More than 3 years

How certain are you about your forecast?

I am very certain about my forecast
I am somewhat certain about my forecast
I am not certain about my forecast

Question 12 When a COVID-19 vaccine becomes available, do you plan to take it?

Definitely not
Probably not
Not sure
Probably yes
Definitely yes

How certain are you about your answer?

I am very certain about my answer
I am somewhat certain about my answer
I am not certain about my answer

Question 13 What is your political leaning?

Definitely Democrat
Usually Democrat
Independent or Other
Usually Republican
Definitely Republican

Question 14 Who will you vote for in the upcoming U.S. Presidential election?

Definitely Joe Biden
Likely Joe Biden
Undecided, Someone else, or Not voting
Likely Donald Trump
Definitely Donald Trump

Question 15 Who do you think will win the upcoming U.S. Presidential election?

Definitely Joe Biden
Likely Joe Biden
Hard to say
Likely Donald Trump
Definitely Donald Trump

How certain are you about your forecast?

I am very certain about my forecast
I am somewhat certain about my forecast
I am not certain about my forecast

Question 16: What is the highest level of school you have completed or the highest degree you have received?

Less than high school degree
High school degree or equivalent (e.g., GED)
Some college but no degree
Associate degree
Bachelor's degree
Graduate degree

Question 17: Which of the following best describes the area you live in?

Urban
Suburban
Rural

Is there any ambiguity in determining your area to be Urban/Suburban/Rural?

Yes
No

Appendix A2: Risk Assessment Survey Questions at the Advisor

Question 1: How would you describe your investment knowledge?

Limited: Virtually no knowledge of financial products
Moderate: Basic knowledge and understanding of financial products and their associated risks
Extensive: Comprehensive knowledge and understanding of financial products and their associated risks

Question 2: What is your expected investment duration?

Less than half a year
Less than 1 year
1 - 3 years
3 - 5 years
Over 5 years

Question 3: How old are you?

Under 25 years old
25 – 35 years
35 – 50 years
Over 50 years old

Question 4: What is your family's annual after-tax income?

Less than 100,000 yuan
100,000 - 300,000 yuan
300,000 - 500,000 yuan
500,000 - 1,000,000 yuan
Over 1,000,000 yuan

Question 5: What is your family's investable assets?

Less than 100,000 yuan
100,000 - 500,000 yuan
500,000 - 2,000,000 yuan
Over 2,000,000 yuan

Question 6: How long have you been investing?

Less than 1 year
1 - 3 years
3 - 5 years
Over 5 years

Question 7: Please select your primary investment category?

Money market funds
Bank wealth management products
Bonds
Stocks

Question 8: What is your investment preference?

4% return with no principal loss
6% return with a maximum 2% potential principal loss
10% return with a maximum 7% potential principal loss
15% return with a maximum 15% potential principal loss
30% return with a maximum 40% potential principal loss

Question 9: What is your primary source of income?

Wages and labor remuneration
Income from business operations
Income from financial assets such as interest, dividends, and asset transfers
Income from non-financial assets such as rental and sale of real estate
No steady income

Question 10: Do you have any large outstanding debts?

No
Yes, long-term debt such as mortgage
Yes, short-term debt such as credit card debts and consumer credit
Yes, loans from relatives and friends

Question 11: Do you have any adverse credit or integrity record?

Yes
No

Online Appendix

1. Probit Analysis

Using the AAII survey data in Section 2, we run the following Probit regression

$$\Pr\left(Answer_{ij} = k\right) = \Phi\left(\alpha + \beta Reverse_i \times Sub_j + \gamma Reverse_i + \delta Sub_j + \varepsilon_{ij}\right), \quad \text{(A1)}$$

for $k = 1, 2, 3, 4, 5$, where $\Phi(\cdot)$ is the CDF of a standard normal distribution. For the case of $k = 1$, the above regression estimates the effect of reversing the response order on the probability of selecting the first choice on the original-order list. Similarly, in the cases of $k = 2$ and $k = 3$, Regression (5) estimates the effects on the probability of selecting the second and third choices. In the cases of $k = 4$ and $k = 5$, Regression (A1) estimates the effect on the probability of selecting the penultimate and last choices.[5] Regression results are reported in Table A1.

The first column of Panel A shows that the coefficient on *Reverse* is -0.00 (*s.e.* = 0.03) and the coefficient on the interaction term is -0.11 (*s.e.* = 0.04). This implies that for objective questions, reversing the response order does not affect the probability of selecting the first choice. For subjective questions, however, the first choice is significantly less likely to be selected when it is presented at the bottom of the list in the reverse-order survey. The implied marginal effects are reported in the first column of Panel B. For subjective questions, reversing the response order reduces the selection probability of the first choice by 1.95 (*s.e.* = 0.44) percentage points, while, for objective questions, the marginal effect of *Reverse* is merely 0.03 (*s.e.* = 0.24) percentage points.

The rest of the table shows that for objective questions, reversing the order has no significant effect on the selection probabilities of all other choices. For subjective questions, however, reversing the order decreases the selection probability by 3.66 (*s.e.* = 0.45) percentage points for the second choice, but increases the selection probability by 1.57 (*s.e.* =0.44) and 3.66 (*s.e.* = 0.42) percentage points for the penultimate and last choices, respectively. For the middle choice ($k = 3$), reversing order appears to have an insignificant effect.

The above results show that respondents have a bias towards the choices at or near the top

[5] Note that the number of choices for a survey question ranges from 2 to 6. In the regressions for the cases of $k = 1$ and $k = 5$, the regression is based on the data from all survey questions. In the regression for the cases of $k = 2$ and $k = 4$, the data from survey questions with 2 or 3 choices are excluded. In the regression for the case of $k = 3$, only the data from questions with 5 choices are included.

of their lists. When the response order is reversed, the last and penultimate choices on the original-order list are now the first and second choices and hence more likely to be selected. The first and second choices on the original-order list are now at the bottom and less likely to be selected. The third choice stays in the middle and reversing order has an insignificant effect. The overall effect is strong for subjective questions and disappears for objective questions to which respondents are likely to have clear answers (such as their age and residence).

2. Decision Time

Our analysis in Section 2 was motivated by the interpretation based on anchoring. For a given survey question, the top choice serves as the anchor and hence respondents tend to select choices close to it, leading to response order bias. For subjective questions, respondents are more uncertain about their answers. Since the anchoring effect tends to be stronger in such circumstances (Jacowitz and Kahneman, 1995), the response-order effect should be stronger. An alternative interpretation of the response order bias is based on the notion of "satisficing behavior" (Krosnick and Presser, 2010). Respondents start from the top of the list and go down until reaching a choice they are satisfied with, leading to an order effect.

Under the satisficing-behavior interpretation, respondents who make decisions more quickly are more likely to be subject to the satisficing behavior; they stop reading the choices as soon as they reach an "acceptable" one and therefore tend to finish their surveys more quickly. Hence, the response-order effect should be stronger among respondents with shorter decision times. The anchoring-based interpretation implies the opposite. If a respondent is more confident about her answers and hence less subject to the order effect, she would also spend less time on her responses. Therefore, the anchoring-based interpretation implies that the order effect should be weaker among respondents with shorter decision times.

To examine these implications, we extend Regression (4) to include a triple interaction term $Reverse_i \times Sub_j \times Fast_i$, where $Fast_i$ is a dummy variable that is 1 if respondent i's survey completion time is among the shortest α ($0<\alpha<1$) of our sample and 0 otherwise. The regression results are reported in Table A2. Column (1) is for the case of $\alpha=1/2$. Consistent with the anchoring-based interpretation, the triple interaction coefficient is -3.45 (*s.e.*=1.15). That is, the order effect is weaker for respondents with shorter decision times. We also repeat the regression with $\alpha=1/3$ and $\alpha=1/4$. The results, reported in Columns (2) and (3), are similar.

3. Quantify the Effect on Survey-based Expectations

In this section, we quantify the response-order effect on survey expectations and contrast it with the effects of two salient features at the time of our survey. The first is the effect of the COVID pandemic. Our survey was conducted from October 27 to November 2, 2020. This was before Pfizer's November 9 announcement of the first major breakthrough in COVID vaccine development. There was substantial uncertainty regarding vaccine development. Xiao and Yan (2023) analyze the data from this AAII survey and show that personal experience of the COVID pandemic plays a disproportionate role in shaping expectations. To measure a respondent's personal experience during the pandemic, we construct two dummy variables: $Mild_i$ is 1 if respondent *i*, according to his answer to Question 8, directly knows someone who has COVID infection with mild or no symptoms but, according to Question 9, does not directly know anyone who has been severely ill due to COVID. $Severe_i$ is 1 if respondent *i*, according to his answer to Question 9, directly knows someone who has been severely ill due to COVID.

The second feature is that our survey was conducted right before the 2020 presidential election. The political division was salient and there was considerable uncertainty about the outcome. A recent literature has shown that an individual's expectations are shaped by her political beliefs (Kempf and Tsoutsoura, 2021; Meeuwis, Parker, Schoar, and Simester, 2022). We construct two dummy variables: R_i is 1 if respondent *i*'s answer to Question 13 on political leaning is "Definitely Republican" or "Usually Republican" and D_i is 1 if respondent *i*'s answer is "Definitely Democrat" or "Usually Democrat." As noted in Table 2, the stated political leanings are subject to the order effect; that is, R_i and D_i are correlated to $Reverse_i$. To isolate the effect from political beliefs, we orthogonalize R_i and D_i against $Reverse_i$ by regressing them on $Reverse_i$ to obtain the residuals, which are denoted R'_i and D'_i, respectively.

To compare the magnitude of the response-order effect with that of the effects of COVID experiences and political leanings, we run the following regression:

$$Y_i = \alpha + \beta_1 Reverse_i + \beta_2 Mild_i + \beta_3 Severe_i + \beta_4 D'_i + \beta_5 R'_i + \varepsilon_i, \qquad \text{(A2)}$$

where Y_i is respondent *i*'s expectation. For ease of interpretation, we set Y_i to be the midpoint value of the respondent's choice range. That is, for the GDP forecast in Question 6, Y_i is -7%, -3.5%, 0%, 3.5%, and 7%, if respondent *i* selected choices 1 through 5, respectively. For the S&P500 return forecast in Question 7, Y_i is -20%, -10%, 0%, 10%, and 20%, if respondent *i* selected choices 1 through 5, respectively. The regression results are reported in Table 3.

In column (1), the coefficient on *Reverse* is 0.95 (*s.e.* = 0.26). That is, reversing the response order increases the S&P500 return forecast by 95 basis points. This is comparable to the effect of COVID experiences. The coefficient on *Mild* is -0.85 (*s.e.* = -0.31), implying that a respondent reduces his return forecast by 85 basis points if he personally knows someone with COVID with mild symptoms. The coefficient on *Severe* is -0.90 (*s.e.* = 0.31). That is, knowing someone with severe illness from COVID reduces a respondent's forecast of the stock market index return by 90 basis points. The order effect is also comparable to the effect of political leanings. The coeffects of D' and R' are -0.63 (*s.e.* = 0.34) and 1.36 (*s.e.* = 0.31), respectively. That is, on average, relative to the forecasts by Independents, Democrats' forecasts are 63 basis points lower while Republicans' are 136 basis points higher. Column (2) shows that the effects on GDP forecast are similar. Reversing the response order increases the GDP growth forecast by 0.45% (*s.e.* = 0.10%). Directly knowing someone who is severely ill from COVID reduces a respondent's forecast by 0.26% (*s.e.* = 0.12%). On average, relative to the forecast by Independents, Democrats' forecast is 0.62% (*s.e.* = 0.13%) lower while Republicans' is 0.86% (*s.e.* = 0.12%) higher.

We then analyze the order effect on respondents' forecast of the vaccine development. Question 11 of the survey is "When do you expect a COVID-19 vaccine to become widely available?" We run Regression (A2) based on the data from this question. For ease of interpretation of regression coefficients, we set the dependent variable Y_i as 1.5, 4.5, 9, 24, and 48 months, if respondent *i* selected choices 1 through 5, respectively. In column (3), the coefficient on *Reverse* is 1.89 (*s.e.*=0.30). That is, when responses are listed in the reverse order, the average forecast for a COVID vaccine is delayed by 1.89 months. The coeffects of D' and R' are 0.78 (*s.e.*=0.39) and -3.61 (*s.e.*=0.36). That is, relative to Independents, Democrats expect a 0.78-month longer wait for a COVID vaccine while Republicans expect a 3.61-month shorter wait.

Finally, we estimate the order effect on respondents' forecast of the election outcome. Questions 15 is about "Who do you think will win the upcoming U.S. Presidential election?" We run Regression (A2) using the responses to this question as the dependent variable. That is, Y_i is 1 through 5 if respondent *i*'s answer is "Definitely Joe Biden," "Probably Joe Biden," "Hard to say," "Probably Donald Trump," and "Definitely Donald Trump," respectively. Column (4) shows that the coefficient on *Reverse* is 0.14 (*s.e.*=0.03), i.e., the reverse-order list (with Trump at the top) leads to a higher predicted chance of a Trump win. The coefficient on *Severe* is -0.10 (*s.e.*=0.04), i.e., respondents who personally know someone seriously ill from COVID are less likely to predict

a Trump win. Interestingly, the coefficients of D' and R' are -0.48 ($s.e.$=0.04) and 0.55 ($s.e.$=0.04), respectively. That is, relative Independents, both Democrats and Republicans predict that their own side has a higher chance to win.

In summary, the above evidence shows that the response order also has significant effects on the expectations on macroeconomic variables but also on variables in other domains such as public health and politics. During our survey period, the COVID-19 pandemic and the presidential election were top of mind and had substantial effects on expectations. Our evidence suggests that the response-order effect is on par with the effect of the dramatic experience during a global pandemic and the effect of political leanings during one of the most divisive presidential elections.

A growing literature uses surveys to measure subjective beliefs (Manski, 2004) to study various topics in belief formation, such as experience effects (Malmendier and Nagel, 2011), partisan biases (Coibion, Gorodnichenko, Weber, 2020), personality traits (Jiang, Peng, and Yan, 2024), and memory (Jiang, Liu, Peng, and Yan, 2025). There has been renewed interest in survey-based beliefs in finance since Greenwood and Shleifer (2014), as they offer unique insights for distinguishing asset pricing theories (see, for example, Choi and Robertson, 2020; Giglio, Maggiori, Stroebel, and Utkus, 2021a, 2021b; De la O and Myers, 2021; Chinco, Hartzmark, and Sussman, 2022; and Bordalo, Gennaioli, LaPorta, and Shleifer, 2025). Our paper adds to the literature by evaluating survey methodologies, highlighting the large response-order bias on survey expectations and preferences.

Our survey is conducted only once and hence our direct evidence is on the *levels* rather than *changes* of the economic variables. However, our evidence does have implications on the potential effect on changes. If a survey is repeated over time, how does the response-order effect influence the changes of a survey-based variable? It is easy to see that if the response order is not consistent across surveys over time, the order effect would make the survey variable more volatile. Suppose the response order stays the same across surveys over time; how does the response-order effect affect the changes in the variables? Let's use the survey of stock return forecast as an example. Suppose we conduct two surveys on the stock return forecast at two points in time. Both have the same list of responses, say, in descending order. The response-order effect implies a positive bias in the forecast in both surveys. When we compute the change in forecast across these two surveys, the two biases would cancel each other out *only if* they have the same size. However, our evidence suggests that the size of the bias is likely increasing in uncertainty. Suppose, for

example, the stock market becomes more uncertain at the time of the second survey. The response-order effect should be stronger, leading to a larger positive bias. Therefore, the change in return forecast is not only biased, but also positively correlated with the market uncertainty, a property that often has important implications to studies in finance.

4. Discontinuity in Self-Directed Investments

The estimated effects in Panel A of Table 6 capture the average treatment effect across all reverse-order users. However, the treatment effect should manifest only when the response order leads to a change in a user's risk category. If the response order alters the user's risk score without crossing a threshold into a different category, no change in investment behavior is expected. In other words, the response-order effect operates primarily through category shifts—when a user's score is pushed across a threshold into a new risk category. Therefore, even a small change in a user's risk score that crosses a category threshold can result in a substantial change in their self-directed investment. This suggests that users with risk scores within a narrow bandwidth around a cutoff should display markedly different investment behaviors depending on whether their scores are above or below the threshold. To test this discontinuity, we focus on users whose risk scores are close to the cutoff between the C4 and C5 categories—the risk score of 69—and estimate the following regression:

$$R5_i = \alpha + \beta 1_{Score_i \geq 69} + \gamma X_i + \varepsilon_i,$$

where $1_{Score_i \geq 69}$ is a dummy variable that is 1 if respondent i's score, $Score_i$, is greater than or equal to 69, and 0 otherwise. X_i denotes a vector of control variables.

Panel A of Table A4 presents the regression results. In column (1), the bandwidth is 3—that is, the regression includes users with risk scores between 66 and 71 (i.e., within 3 points of the cutoff at 69). The estimate of β is 1.42 (*s.e.* = 3.08), indicating that users whose scores are just above the threshold (and are therefore classified as C5) invest RMB 1420 more R5 products than users just below the cutoff. Columns (2) and (3) narrow the bandwidth to 2 and 1, respectively, and the estimates of β remain qualitatively similar. In summary, there is a substantial discontinuity in the investment in R5 products at the score threshold of 69, the boundary between risk categories C4 and C5.

We conduct two placebo tests. Panel B reports the results using a fictitious cutoff point of 66. In this case, all users in the sample fall within the C4 category. As expected, the estimates of

β are close to zero with small standard errors, indicating that users just above and below this artificial threshold allocate similar amounts to R5 products. Likewise, Panel C presents results based on a fictitious cutoff point of 72, where all users are in the C5 category. Again, the estimates show no meaningful difference in the investment in R5 products across the artificial threshold, consistent with the absence of a risk category change.

We note that the estimated effect in Panel A should not be interpreted as causal. In principle, users on either side of the cutoff may have self-selected into their positions, and such selection could partly explain the large difference in portfolio allocations. Nonetheless, the evidence lends support to the compliance mechanism by demonstrating a strong association between risk category and investment behavior—even when the underlying risk score changes only marginally.

References

Bordalo, Pedro, Nicola Gennaioli, Rafael LaPorta, and Andrei Shleifer, 2025, Finance Without Exotic Risk, *Journal of Financial Economics*, forthcoming.

Choi, James and Adriana Z. Robertson, 2020, What Matters to Individual Investors? Evidence from the Horse's Mouth, *Journal of Finance* 75, 1965–2020.

Chinco, Alex, Samuel Hartzmark, and Abigail Sussman, 2022, A New Test of Risk Factor Relevance, *Journal of Finance* 77, 2183–2238.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber, 2020, Political Polarization and Expected Economic Outcomes, Working paper.

De la O, Ricardo, and Sean Myers, 2021, Subjective cash flow and discount rate expectations. *Journal of Finance* 76, 1339–1387.

Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of Returns and Expected Returns, *Review of Financial Studies* 27, 714–746.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2021, Five facts about beliefs and portfolios, *American Economic Review* 111, 1481–1522.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2021, The joint dynamics of investor beliefs and trading during the COVID-19 crash, *PNAS* 118 (4) e2010316118.

Jacowitz, Karen E. and Daniel Kahneman, 1995, Measures of Anchoring in Estimation Tasks, *Personality and Social Psychology Bulletin* 21, 1161–1166.

Jiang, Zhengyang, Hongqi Liu, Cameron Peng and Hongjun Yan, 2025, Investor Memory and Biased Beliefs: Evidence from the Field, *Quarterly Journal of Economics*, forthcoming.

Jiang, Zhengyang, Cameron Peng and Hongjun Yan, 2024, Personality Differences and Investment Decision-Making, *Journal of Financial Economics* 153, 103776.

Kempf, Elisabeth and Margarita Tsoutsoura, 2021, Partisan Professionals: Evidence from Credit Rating Analysts, *Journal of Finance* 76, 2805–2856.

Krosnick, Jon and Presser, Stanley, 2010, Question and Questionnaire Design, in *Handbook of Survey Research* (2nd Edition) James D. Wright and Peter V. Marsden (Eds). San Diego, CA: Elsevier.

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression babies: Do Macroeconomic Experiences Affect Risk-taking? *Quarterly Journal of Economics* 126, 373–416.

Manski, Charles, 2004, Measuring Expectations, *Econometrica*, 72, 1329–1376.

Meeuwis, Maarten, Jonathan A. Parker, Antoinette Schoar, and Duncan I. Simester, 2022, Belief Disagreement and Portfolio Choice, *Journal of Finance* 77, 3191–3247.

Xiao, Yaqing and Hongjun Yan, 2023, Relatable Information Disproportionately Shapes Expectations, Working paper.

Table A1. Probit Regressions

This table is based on the AAII survey data. Panel A reports the results of the following Probit regression for k=1, 2, 3, 4, 5:

$$\Pr\left(Choice_{ij} = n\right) = \Phi\left(\alpha + \beta Reverse_i \times Sub_j + \delta Reverse_i + \gamma Sub_j + \varepsilon_{ij}\right),$$

where $Choice_{ij}$ is respondent i's choice for Question j; k=1, 2, 3, 4, 5 refer to the first, second, third, penultimate, and last choices, respectively; $\Phi(\cdot)$ is the CDF of a standard normal distribution; $Reverse_i$ is a dummy variable that is 1 if respondent i faces a reverse-order survey and 0 otherwise; and Sub_j is a dummy variable that is 1 if Question j is a subjective question and 0 otherwise. The classification of subjective and objective questions is reported in Table 1. For k=2 and k=4, the data from survey questions with 2 or 3 responses are excluded. For k=3, the regression is based on the data from only the Questions with 5 responses. Survey-question fixed effects are included in all specifications. Panel B reports the implied marginal effect of *Reverse*. Standard errors, in parentheses, are based on standard errors that are clustered by respondent. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

Panel A. Probit regressions

	First (k=1)	Second (k=2)	Middle (k=3)	Penultimate (k=4)	Last (k=5)
Reverse	-0.00	-0.08	-0.03	-0.02	0.04
	(0.03)	(0.07)	(0.05)	(0.03)	(0.02)
Reverse × *Sub*	-0.11***	-0.07	-0.04	0.08**	0.13***
	(0.04)	(0.07)	(0.05)	(0.03)	(0.03)
Sub	1.33***	2.05***	0.81***	-0.82	-1.83***
	(0.10)	(0.07)	(0.05)	(0.04)	(0.05)
Constant	-2.64***	-2.04***	-1.38***	-0.45***	0.16***
	(0.10)	(-33.06)	(0.04	(0.03)	(0.03)
Question FEs	Yes	Yes	Yes	Yes	Yes
N	49779	40454	34195	40454	49779
Pseudo R-squared	0.352	0.190	0.162	0.074	0.297

Panel B. Marginal effects of *Reverse* (%)

	First (k=1)	Second (k=2)	Middle (k=3)	Penultimate (k=4)	Last (k=5)
Subjective questions	-1.95***	-3.66***	-1.78***	1.57***	3.66***
	(0.44)	(0.45)	(0.55)	(0.44)	(0.42)
Objective questions	-0.03	-0.12	-3.89	-0.71	0.88
	(0.24)	(0.10)	(0.69)	(1.07)	(0.56)

Table A2. Variations across Decision Time

This table reports the results from the regression that expands the regression in Panel A of Table 2 with a triple interaction term $Reverse_i \times Sub_j \times Fast_i$, where $Fast_i$ is a dummy variable that is 1 if respondent i's survey completion time is among the shortest α (=1/2, 1/3, or 1/4) of the respondents' completion times. Survey-question fixed effects are included in all specifications. T-statistics, in parentheses, are based on standard errors that are clustered by respondent. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

Dep. var. = $Score$	(1) α=1/2	(2) α=1/3	(3) α=1/4
Reverse × *Sub* × *Fast*	-3.45***	-3.48***	-2.59**
	(1.15)	(1.22)	(1.32)
Reverse × *Sub*	5.48***	4.88***	4.37***
	(0.81)	(0.71)	(0.67)
Sub × *Fast*	0.74	0.76	0.18
	(0.84)	(0.87)	(0.94)
Reverse × *Fast*	2.91***	3.17***	2.11*
	(0.93)	(0.99)	(1.09)
Reverse	-1.10	-0.68	-0.15
	(0.65)	(0.57)	(0.54)
Sub	-12.73***	-12.62***	-12.41***
	(0.81)	(0.76)	(0.74)
Fast	-0.81	-0.86	-0.41
	(0.67)	(0.71)	(0.78)
Constant	50.21***	50.10***	49.91***
	(0.64)	(0.60)	(0.58)
Question FEs	Yes	Yes	Yes
N	52908	52908	52908
R-squared	0.318	0.318	0.318

Table A3. Quantify the Order Effect on Expectations

This table reports the results from the following regression:

$$Y_i = \alpha + \beta_1 Reverse_i + \beta_2 Mild_i + \beta_3 Severe_i + \beta_4 D_i' + \beta_5 R_i' + \varepsilon_i,$$

where Y_i is respondent i's answer to a survey question; $Reverse_i$ is a dummy variable that is 1 if respondent i's survey is in the reverse order and 0 otherwise; $Mild_i$ is a dummy variable that is 1 if respondent i answered "Yes" to Question 9 ("Do you directly know someone who has been infected by COVID-19 with mild or no symptoms?") and "No" to Question 10 ("Do you directly know someone who has been severely ill due to COVID-19 infection?"); $Severe_i$ is a dummy variable that is 1 if respondent i answered "Yes" to Question 10; D_i' is the residual from regressing D_i on $Reverse_i$; R_i' is the residual from regressing R_i on $Reverse_i$; D_i is a dummy variable that is 1 if respondent i answered "Definitely Democrat" or "Usually Democrat" to Question 13 ("What is your political leaning?"); and R_i is a dummy variable that is 1 if respondent i answered "Definitely Republican" or "Usually Republican" to Question 13. In column (1), Y_i is the forecast of the S&P 500 stock index return and is set to -20%, -10%, 0%, 10%, and 20% if respondent i selected choices 1 through 5, respectively. In column (2), Y_i is the forecast of the GDP growth rate and is set to -7%, -3.5%, 0%, 3.5%, and 7% if respondent i selected choices 1 through 5, respectively. For column (3), the survey question is "When do you expect a COVID-19 vaccine to become widely available?" For these two regressions, Y_i is 1.5, 4.5, 9, 24, and 48 months if respondent i selected choices 1 through 5, respectively. For column (4), the survey question is "Who do you think will win the upcoming U.S. Presidential election?" Y_i is 1 through 5 if respondent i's answer is "Definitely Joe Biden," "Probably Joe Biden," "Hard to say," "Probably Donald Trump," and "Definitely Donald Trump," respectively. All regressions include age, gender, education, and residence fixed effects. Standard errors are reported in parentheses. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

	Stock return (%)	GDP growth (%)	Vaccine dev. (m)	Election outcome
	(1)	(2)	(3)	(4)
Reverse	0.95***	0.45***	1.89***	0.14***
	(0.26)	(0.10)	(0.30)	(0.03)
Mild	-0.85***	-0.13	-0.63*	-0.002
	(0.31)	(0.12)	(0.35)	(0.04)
Severe	-0.90***	-0.26**	-0.05	-0.10**
	(0.31)	(0.12)	(0.36)	(0.04)
D'	-0.63*	-0.62***	0.78**	-0.48***
	(0.34)	(0.13)	(0.39)	(0.04)
R'	1.36***	0.86***	-3.61***	0.55***
	(0.31)	(0.12)	(0.36)	(0.04)
Constant	-1.60	-1.97*	18.70***	3.40***
	(2.91)	(1.15)	(3.36)	(0.37)
N	3,009	3,027	3,026	3,007
R-squared	0.035	0.071	0.091	0.226

Table A4. Discontinuity in Self-Directed Investments

This table, based on the data from the RCT in Section 3, shows the discrete jump in investors' purchases of R5 products at the risk score 69. Specifically, we run the following regression:

$$R5_i = \alpha + \beta 1_{Score_i \geq K} + X_i + \varepsilon_i,$$

where $R5_i$ is user i's purchase of R5 products and is denominated in RMB 1000, $1_{Score_i \geq K}$ is a dummy variable that is 1 if respondent i's risk assessment score, $Score_i$, is greater than or equal to K, and 0 otherwise. Control variables includes *Male, Age, LnInitialDirect, NewUser*, all of which are defined in Table 3. The regressions are based on the users whose scores are within n points from the cutoff point K. For example, in column "$n = 3$", investors' final scores are within 3 points of the cutoff point K (i.e., $K - 3 \leq Score_i \leq K + 2$). The table only reports the estimated coefficients on $1_{Score_i \geq K}$. $K = 69,\ 66$, and 72 in Panels A, B, and C, respectively. Robust standard errors are reported in parentheses. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

Panel A. $K = 69$

	$R5$		
	$n = 3$	$n = 2$	$n = 1$
$1_{Score_i \geq K}$	1.42***	1.60***	0.84*
	(0.35)	(0.56)	(0.49)
Controls	Y	Y	Y
N	1,525	950	505
R-squared	0.10	0.09	0.11

Panel B. $K = 66$

	$R5$		
	$n = 3$	$n = 2$	$n = 1$
$1_{Score_i \geq K}$	0.14	0.07	-0.04
	(0.11)	(0.06)	(0.13)
Controls	Y	Y	Y
N	1,193	832	402
R-squared	0.00	0.02	0.04

Panel C. $K = 72$

	$R5$		
	$n = 3$	$n = 2$	$n = 1$
$1_{Score_i \geq K}$	0.00	-0.21	-0.06
	(0.48)	(0.57)	(0.77)
Controls	Y	Y	Y
N	1,534	1,157	500
R-squared	0.17	0.18	0.24