The Demand for Robo-Advising: Adoption and Delegation Intensity

Richard Harris* Chaojie (Jay) Liu† Manuela Pedio‡

July 30, 2026

This draft is being updated regularly. For the latest version, please check: here

Abstract

Technology has made financial advice cheaper and easier to reach than it was, yet the households who would gain most from it still use it least. We ask what stands in the way, using monthly proprietary data on 100,000 Chinese mutual fund investors and a robo-advisor that executes the portfolios it recommends, so that adoption and the share of fund wealth delegated are observed separately. Underdiversification and total turnover lower adoption but raise the share adopters delegate; confirmation bias raises adoption and lowers that share. Experienced fund investors are the slowest to adopt, and those who do adopt delegate less than others. Three explanations fit. Investors correct the mistakes delegation makes visible, fastest among those whose portfolios were the most concentrated or most heavily traded to begin with and who follow the platform most closely, which we call selective awareness. They surrender fund selection reluctantly even where they accept trading automation. And experienced adopters reclaim control as delegation tenure accumulates.

Keywords: Household Finance, Financial Advice, Robo-Advising, Behavioral Finance, FinTech

JEL Classification: G11, G41, D83, G24, O33

1. Introduction

Households make investment mistakes that are large, persistent, and concentrated among those least equipped to avoid them (Calvet et al., 2007; Campbell, 2006). Underdiversification and excessive trading alone cost retail investors a material share of their returns (Barber & Odean, 2000; Goetzmann & Kumar, 2008). On the available evidence, traditional human financial advice has reached these households poorly. Advisory fees and account minimums price out many of the small accounts where mistakes concentrate, and advisers have been found to recommend what they themselves wrongly believe, or what their incentives favor, rather than what the client needs (Chalmers & Reuter, 2020;

*University of Bristol Business School, Richard.Harris@bristol.ac.uk.

†University of Bristol Business School, Jay.Liu@bristol.ac.uk.

‡University of Bristol Business School and Baffi CAREFIN Centre, Manuela.Pedio@bristol.ac.uk. The authors acknowledge and appreciate the support from the Digital Economy Open Research Platform (www.deor.org.cn). All data were sampled and desensitized, and were remotely analyzed on the Ant Open Research Laboratory in an Ant Group environment. We thank Chris Brooks, Miklos Farkas, Lorenz Kueng, Ian Tonks, Guanlin Wang, and brown bag participants at USI Lugano for useful feedback. Jay particularly thanks Jian Hou and Wenhao Zhang from Ant Group for their help with data processing and technical issues. The opinions in this paper represent only those of the authors and do not necessarily represent those of Alipay or Ant Group. All mistakes are our own.

Egan et al., 2019; Foerster et al., 2017; Hackethal et al., 2012; Linnainmaa et al., 2021). Robo-advisors are the leading response. They build and rebalance diversified, rule-based portfolios for fees of a few tenths of a percent, they reach account sizes human advice rarely reached (Reher & Sokolinski, 2024), and among the investors studied so far they improve diversification and reduce trading mistakes (Bianchi & Brière, 2026; D'Acunto et al., 2019; Hao et al., 2022; Reher & Sokolinski, 2024; Rossi & Utkus, 2024); for surveys of what the technology promises and what it has so far delivered, see D'Acunto and Rossi (2023) and, on financial inclusion more broadly, Philippon (2019).

Barriers to entry and conflicted incentives are unlikely to be the whole story, however. The sharpest evidence we are aware of comes from Bhattacharya et al. (2012), whose field experiment removed both at once, offering advice that was free and unbiased to roughly 8,000 retail clients of a German brokerage. About five percent obtained it, the investors who by the authors' need measures stood to gain most were the least likely to obtain it, and the few who took it hardly followed it. Whatever stands between households and financial advice, then, is not price or access alone. What else stands there is what we do not know, and the question matters because the welfare gains documented in the effects literature are all conditional on use.

This paper therefore asks what still prevents investors from taking up financial advice once price and access no longer do. The question cannot be answered on its own, because adoption, the decision to enter, is not the only decision a household makes. Having adopted, it must also decide how much of its wealth to place under the algorithm, and having decided that, whether to keep it there. Those two decisions are what we mean by commitment, and both are directly observable: the delegated share, the fraction of its fund wealth it places with the algorithm rather than in funds it picks itself, and the path of that share as tenure lengthens. We therefore ask what governs the demand for delegated advice at entry and in commitment, and in our data the characteristics that predict entry are not the characteristics that predict commitment. The investors whose portfolios most need correcting are the slowest and least likely to adopt, but they delegate the most once they have tried it. For the same household, what deters entry and what limits commitment are different economic objects.

Entry and commitment ask different things of the investor. Entry requires discovery, evaluation, and enough initial trust to try the product. Commitment requires her to hand over the choice of which funds to hold and trade, and it is visible in how much of her fund wealth she moves across and how long she leaves it there. A concentrated investor may be slow to adopt because she does not perceive concentration as costly, yet delegate heavily once she adopts and sees a diversified algorithmic portfolio. An experienced fund investor may know enough to evaluate the product quickly, yet delegate little because fund selection, which is automated by robo-advising, is the decision she values. A binary use indicator cannot tell these apart, which is one reason the literature does not agree on who adopts robo-advice (Baulkaran & Jain, 2023; Belanche et al., 2019; Brenner & Meyll, 2020; D'Acunto et al., 2019; D'Hondt et al., 2020; Piehlmaier, 2022; Scherer & Lehner, 2025).

Testing this requires evidence rarely available in one setting, and the constraint is one the robo-advising literature has named from the inside. D'Acunto et al. (2019) report that the biases they study "were not fully eliminated, because investors could place additional trades on top of what the robo-adviser suggested," so how much an investor hands over bounds what adoption can accomplish, and they do not observe it. Surveys measure willingness to use advice but not realized delegated

shares. Robo-advising data record product use but often not the self-directed fund holdings and trades against which delegation has to be measured. Portfolio datasets observe holdings but rarely allow a comparison between delegated and self-directed investments inside the same account. The closest existing setting is the retirement plan, where participants who hold a target-date fund alongside other funds do reveal a delegated share (Mitchell & Utkus, 2022), but not one shaped by demand in the way we need.[1]

We study BangNiTou, a fully delegated mutual fund robo-advisor founded by Ant Group and Vanguard and embedded in Alipay, with asset-based fees of 0.20 to 0.50 percent. The same account records self-directed fund holdings, delegated robo assets, automatic investment plans, cash-management balances, platform engagement, transactions, and realized returns. What makes the setting useful is that it records three ways of holding funds that differ in how much of the decision the investor hands to algorithms. She can pick and trade funds herself, which we label *DIRECT*. She can set an automatic investment plan that buys the funds she chose on a schedule, which we label *AUTO*, and which automates the timing of her purchases while leaving the choice of what to own with her. Or she can hand both fund selection and trading to the algorithm, which is the robo-advising we study, *RA*. Every investor in the sample already delegates security selection to fund managers, since all of them invest through mutual funds, so comfort with delegated management as such is held fixed. The decision rights the three products divide are the choice of which funds to hold and the trading of them.

We observe 100,000 mutual fund investors monthly from September 2020 through September 2024, 4.9 million investor-months, and we build ten behavioral measures from each investor's own fund trades before the decisions we model. Of these investors, 70,034 use RA at least once and 29,966 hold mutual funds but never use it. The sample is deliberately stratified on ever adoption of RA, so its base rates are not platform penetration.[2]

Because the delegated share is observed only for adopters, we estimate Heckman selection models in which whether an investor adopts and how much she delegates once she has are modeled jointly, with platform tenure excluded from the share equation as a discovery shifter. We report unweighted estimates, since the object of interest is a slope within the sampled population rather than a population prevalence; the weighted re-estimation in the appendix leaves the cross-margin sign pattern in place. We estimate them month by month and report Fama–MacBeth averages with Newey–West standard errors (Fama & MacBeth, 1973; Newey & West, 1987), because the characteristics of interest move slowly or not at all and investor fixed effects would absorb the cross-sectional variation the question depends on. The exclusion restriction is the main identifying assumption, and our estimates should be read as conditional on it; Section 7.5 reports what happens when it is relaxed.

The estimates show that entry and commitment, the two margins, are not governed by the same characteristics, and that several characteristics take opposite signs across them. Wealth is almost entirely an entry variable: a one-standard-deviation increase in log total financial assets raises adoption

[1] In a retirement plan the menu is set by the employer, entry is frequently a default rather than a choice, and the non-delegated sleeve is drawn from the same restricted menu instead of being built by the household. Mitchell and Utkus (2022) report that mixed target-date investors, those directing part of their contributions to a target-date fund, are roughly a third of target-date holders, so the margin exists in the data; it is the demand interpretation that does not carry over.

[2] The stratification buys power on rare events at the cost of the unconditional prevalence, which the design cannot recover (Manski & Lerman, 1977). The estimand is the cross-sectional slope within the sampled strata, and every magnitude below is read against a sample base rate; Section 4.2 gives the detail.

by 14.2 percentage points, about 44 percent of the participation mean, while its negative relation to the delegated share falls to near zero once behavioral controls enter. Fund experience raises pooled participation but lowers the conditional share, and it lowers first adoption outright. Risk preference raises both margins, which is what a characteristic that does not separate the margins looks like.

The behavioral measures generate the paper's central pattern. Concentration is the Herfindahl index of an investor's mutual fund weights, which is high when she holds few funds or spreads her wealth unevenly across them. It is the measure this literature most often uses for the mistake a diversified rule-based portfolio is built to correct, and the appendix reports the complementary diversification measures, the effective number of holdings and the weighted return correlation, alongside it. A one-standard-deviation increase in portfolio concentration lowers adoption by 6.5 percentage points, about 17 percent of the 37.2 percent participation mean, and raises the delegated share among adopters by 2.0 percentage points, about 13 percent of the 15.1 percent conditional mean. Turnover behaves the same way, lowering adoption by 3.0 percentage points and raising the delegated share by 4.5, which is 8 percent of the adoption rate and 30 percent of the average delegated share. Confirmation bias runs the other way, raising adoption by 3.0 percentage points and lowering the delegated share by 2.2. Attention, gambling preference, and herding lower both margins. The investors who make the mistakes the product is built to correct are therefore the slowest to adopt it and the deepest delegators once they have. The pattern is not confined to the pool of ongoing users: at the investor's very first delegation decision, concentration lowers first adoption and raises the initial delegated share by 4.5 percentage points, so it appears at the moment of entry and not only among those who stay.

Opposite signs establish that a single use indicator is too coarse for a delegated product, because it averages entry and commitment, which move apart. They do not identify why they move apart. More than one account generates the same signs, and the cross-section that produced them cannot choose among those accounts. Each of the three mechanisms below is therefore tested on evidence the baseline does not contain, and a fourth possibility, that recent returns drive all of it, is tested against the same standard.

The first mechanism is selective awareness. If exposure to a diversified, rule-based portfolio is what makes the cost of a concentrated one visible, then the mistake should recede after adoption, recede most among the investors who were worst at the start, and recede faster the more the investor looks at the platform. All three hold. Concentration and turnover decline with months of use; the declines are far larger for adopters in the top decile of each mistake before adoption, an additional 0.13 of concentration and 0.19 of turnover over a year of use, which is 0.4 and 0.5 of each measure's pooled standard deviation; and the decline steepens with login frequency for both. None of the comparison biases shows that engagement dose, which is the discriminating test, because mean reversion or a maturing portfolio would not produce a steeper decline for investors who look more often. Login frequency is neither randomly assigned nor specific to the robo-advisor, so it may also stand for sophistication or general attentiveness; what the test uses is that those alternatives should dose every bias alike, and they do not. Awareness is selective across mistakes rather than uniform, which is also why attention, gambling, and herding depress both margins: delegation does not relax the constraint that generates them. The reading has a rational benchmark, since investors who optimally learn about the assets they already hold can find concentration optimal rather than mistaken, so its cost need never

become salient to them (Van Nieuwerburgh & Veldkamp, 2010), concentrated holdings can also be consistent with rational preferences rather than error (Polkovnichenko, 2005), and portfolio attention is selective, not continuous (Sicherman et al., 2016).

The second mechanism is decision authority. The robo-advisor does two things at once: it automates buying, and it takes over the choice of which funds to hold. An investor who stays away from it may be avoiding the automation, or may be avoiding giving up that choice. The automatic plan separates the two, since it automates the timing of her purchases while she still picks the funds herself. We stack each investor-month into its DIRECT, AUTO, and RA rows with investor-month fixed effects, which difference out wealth, market conditions, and any common taste for automation, leaving only how the same household in the same month uses the three products relative to one another. The coefficients are incremental: one gives the step from self-direction to automation, the other the further step from automation to full delegation. The characteristics that move the first step are not the ones that move the second, and several move them in opposite directions. Concentration raises the automation step by 5.8 percentage points and lowers the delegation step by 3.3, and turnover moves the two steps in the same directions, at 3.5 and -3.2; engagement and attention bias split in the opposite direction, lowering the automation step by 4.2 percentage points each and raising the delegation step by 1.6 and 2.6. A common taste for automation implies no differential at the second step, since a product that automates more would attract such an investor no less than one that automates less. The second step is instead significantly non-zero and opposite in sign to the first for each of these four characteristics and for anchoring, which is what a common taste cannot produce. We build here on experimental evidence that decision rights carry intrinsic value beyond their instrumental payoff (Bartling et al., 2014) and that people retain authority even when retention is costly (Fehr et al., 2013), and we take that primitive to a portfolio setting where the value of the right shows up in the choices investors actually make rather than in what they say. The design also separates what the algorithm-aversion experiments bundle: Dietvorst et al. (2018) restore algorithm use by letting subjects modify its output, and AUTO is the field counterpart of that retained control.

The third mechanism is reclaiming control. It concerns commitment over time rather than at a point, and it explains a baseline result the first two do not: fund experience raises pooled participation while lowering the conditional share, and lowers first adoption outright. If experience raises the value of keeping fund selection, adoption should not be the end of the matter. Among active adopters, longer RA tenure goes with lower delegated shares, and the gradient is steeper for investors with more pre-adoption fund experience, at -0.023 to -0.043 on the interaction. The self-directed side shows the same pattern: portfolio concentration falls sharply at adoption and rebuilds toward its pre-adoption level over roughly two years. Delegation is not an absorbing state, and for experienced investors it can be partial or transitional.

Recent performance is the leading alternative to all three, and it narrows the account rather than displacing it. If households simply chase returns, the characteristic gradients would proxy for performance feedback. Performance does enter demand, but in a structured way: through the magnitude of own returns at first adoption, where gains and losses both deter entry, and through their sign at exit, where good self-directed returns push investors out. It absorbs none of the characteristic gradients. A residual channel also remains. Selection on unobservables is positive when the sample is left

unweighted, and it does not weaken as controls are added: the coefficient on the selection correction rises from 0.158 to 0.212 once skills and the behavioral measures enter, so the marginal adopter drawn in by something we do not observe delegates more, not less. The term is the one coefficient in the paper that does not survive population reweighting, and we read it accordingly. Two features of that estimate point to what the unobservable is. Its sign rules out the readings that would otherwise be natural, since an investor who enters tentatively, or who is adversely selected into the product, would delegate less than average and produce a negative coefficient. Its stability rules out the observables as substitutes, because measured skill and ten behavioral biases fail to absorb it. What survives both is a disposition that raises entry and commitment together, which is how the money-doctor theories describe trust in delegated management (Gennaioli et al., 2015; Guiso et al., 2008). We therefore offer trust as an interpretation of the residual rather than a finding, noting that our design identifies it only through the covariance between the two equations, not as a measured belief. The evidence throughout is observational: the mechanism labels rest on sign patterns, timing, and dose-response, not on randomized variation in beliefs or in how much authority the investor gives up, and we do not estimate causal effects of adoption.

One implication follows, and it is a hypothesis for a platform to test rather than a result we establish. Because selective awareness appears to gate entry, a platform that wants to reach underdiversified or heavily trading investors should make their own mistakes visible to them before they decide, and that is the intervention to run against advertising the algorithm's past performance.

The paper proceeds as follows. Section 2 places the paper in the literature and states what it contributes. Section 3 sets out the two-margin framework and shows why opposite signs are a generic outcome when entry costs and benefit loadings differ. Section 4 describes the setting, the data, and the designs. Section 5 reports the baseline Heckman evidence, the first-adoption evidence, and why the baseline requires mechanism tests. Section 6 tests the three mechanisms and the performance alternative. Section 7 reports post-adoption consequences and robustness. Section 8 concludes.

2. Related Literature and Contribution

We make four contributions. First, separating adoption from conditional delegation reconciles the contrasting findings on who uses robo-advice (Baulkaran & Jain, 2023; Belanche et al., 2019; Brenner & Meyll, 2020; D'Acunto et al., 2019; D'Hondt et al., 2020; Piehlmaier, 2022; Scherer & Lehner, 2025). A binary indicator pools entry with commitment, so samples that differ in composition recover different mixtures of two margins on which the same characteristic can take opposite signs, and prior results can disagree without any of them being wrong.

The reconciliation is sharpest against Piehlmaier (2022), who finds that overconfident investors adopt more, and it turns on being precise about which overconfidence is at stake. His measure is survey-based, self-assessed confidence in excess of tested knowledge. Ours are revealed from trading, where the literature offers several proxies; we report the two that bear on delegation, total turnover (Barber & Odean, 2000, 2001) and speculative turnover (Barber & Odean, 2002; Han & Kumar, 2013), and enter them together. Total turnover lowers adoption while raising the delegated share, so a use indicator recovers a different mixture of the two margins in every sample, and the two findings

are compatible once entry is separated from commitment. Speculative turnover, held apart from total turnover, lowers both margins instead, and the margin separation is not what reconciles that one. It records the belief that one knows when to be out of a position and when to be back in, which a rule-based portfolio cannot act on, so these investors neither enter nor commit.[3]

The same separation adds a step to a distinction household finance already draws. That literature separates whether a household participates in a market from how well it invests once it does (Calvet et al., 2007; Campbell, 2006). For a product the household does not run itself, those two are no longer adjacent. Adopting a robo-advisor does not determine the portfolio; the investor still chooses what fraction of her fund wealth to place under it, and she keeps managing the rest. Her portfolio quality is a weighted average of the two sleeves, the part the algorithm runs and the part she still runs herself, with the delegated share as the weight, so the number of adopters and the quality of the portfolios they end up holding can move apart. A policy that raised adoption while leaving delegated shares where they are would improve portfolios far less than the adoption figure suggests.

Second, we treat behavioral investment patterns as the object governing demand for a corrective technology, not as nuisance controls. Behavioral finance documents costly household mistakes (Barber & Odean, 2000, 2001, 2008; Goetzmann & Kumar, 2008; Kumar, 2009), and the effects literature shows algorithms can reduce them conditional on use (Bianchi & Brière, 2026; D'Acunto et al., 2019; Reher & Sokolinski, 2024; Rossi & Utkus, 2024). Both take use as given. We supply the selection step in between, which matters because the aggregate gains from a corrective technology depend on whether the investors who would gain most expose themselves to it. Rossi and Utkus (2024) come closest, reporting that the investors who benefit most from a hybrid robo-advisor are also the most likely to adopt and to remain enrolled. Our finding runs the other way, and the difference is one of object rather than of result. Their benefit is a gain the investor has already realized, measured among people who had enrolled, so their selection statement is conditional on adoption and speaks to who stays. Our need is built from the investor's own trades in the months before she decides, while she is still a non-adopter, so ours is unconditional and speaks to who enters at all. Both can hold at the same time, since the investors who would gain most may be slow to arrive and, once they arrive, the right ones may stay.[4]

Third, we bring the economics of decision rights to household portfolio choice, and we measure what it costs a household to give one up. Experimental work establishes that a decision right carries value beyond the payoff it delivers, that people keep authority even when keeping it is costly, and that holding a right and exercising it are separable (Aghion & Tirole, 1997; Bartling et al., 2014; Fehr et al., 2013). That work reads the value of authority off willingness to pay in the laboratory, where the right is an abstract one and the price of keeping it is set by the experimenter. Household finance has not made much use of the idea, in part because portfolio data rarely contain a moment at which a household visibly hands a decision over. This setting contains one, and the right at stake is concrete: who chooses which funds to hold and places the trades that implement the choice.

[3] Stated confidence and realized overtrading need not move together, and the two revealed measures need not either. Which proxy a study uses is therefore itself part of why the literature disagrees.

[4] The settings also differ in ways that would push the same way. Theirs is a hybrid advisor retaining a human adviser and serving prior self-directed clients with far larger accounts; ours is a mass-market platform with small accounts and effectively no minimum.

The measurement follows from the product menu rather than from an elicited valuation. The same investor in the same month can direct her own portfolio, buy an automatic plan that schedules her purchases while she keeps fund selection, or delegate fund selection to the robo-advisor. Comparing the automatic plan with the robo-advisor inside the same investor-month holds the platform, the interface, and the mechanical work of trading fixed, so what is left between the two products is who chooses the funds. The two steps do not move together: portfolio concentration raises the automation step by 5.8 percentage points while lowering the delegation step by 3.3, and total turnover moves the two steps in the same directions. Investors who hold active views about which funds to own tilt toward the product that automates the work of trading and away from the one that takes the decision, which is what an authority cost predicts and what a common taste for automation cannot produce. The two coefficients are incremental, so this is a relative tilt: concentrated investors still use RA more than they self-direct, by 2.5 percentage points net, but they favor the fund-choice-preserving product by more. Keeping the choice of funds therefore costs the investor something in this market, and she pays it in delegated wealth rather than in a stated willingness to pay.

That reading also speaks to the algorithm-aversion literature, where subjects abandon an algorithm after seeing it err (Dietvorst et al., 2015) and return to it when allowed to adjust its output (Dietvorst et al., 2018), while other work finds that people prefer algorithmic to human judgment in the first place (Logg et al., 2019). Those designs necessarily bundle two things, accepting the algorithm and surrendering the decision, because the subject has no way to accept one without the other. Where an experiment does let the two come apart, it finds the same thing we do about what investors are actually resisting: Holzmeister et al. (2023) report that subjects would rather delegate to an algorithm than to a compensated professional, so resistance is not to the machine, and that subjects who rate their own decision-making more highly delegate less. The second half is the authority cost, and it is the pattern our fund-experience result takes. The product menu here separates them, and that separation is what our design uses: what investors resist is not the algorithm, since they buy the automated plan, but giving up the choice of funds. The distinction carries beyond robo-advice to managed accounts, target-date funds, and advisor-sold products, which automate different amounts of the work and take different amounts of the decision.

Fourth, we add a question to the literature on the demand for financial advice: not only whether a household accepts advice, but how much of its portfolio it hands over once it has. That literature asks whether households obtain advice when it is available and affordable, and finds that many do not, that the ones who would gain most are often the least likely to, and that advice is frequently miscalibrated or conflicted when it does arrive (Bhattacharya et al., 2012; Foerster et al., 2017; Hackethal et al., 2012; Linnainmaa et al., 2021; von Gaudecker, 2015). The work that does follow households after they say yes has looked at what the adviser does rather than at how much the household keeps for itself: Greig et al. (2025) study a hybrid robo-advisor in which portfolio management is automated and human advisers supply everything else, and trace retention to what those advisers do. We hold the adviser fixed, since ours is an algorithm, and let the household choose how much of the decision it hands over. The closest paper is Bhattacharya et al. (2012), the experiment that opens this one, and the first thing to say about it is that it asks a different question. They ask whether supplying good advice is sufficient, and answer it by supplying some: the treatment is the offer, the outcome is whether clients

accept and follow it, and the finding is that most do not. We take that answer as our premise and ask what the household weighs once the offer is not the binding constraint. Their question is about the supply of advice; ours is about the demand for it. The two share a subject and a puzzle, so the comparison is worth drawing. Both concern retail investors who would gain from correcting their portfolios, both find that the investors with the most to gain are the least likely to obtain help, and both treat that gap as the thing to be explained rather than as noise.

Three features of our setting are what let us take the demand question further. Their advice was a recommendation the client still had to carry out herself, so a client who ignored it and a client who never wanted it look the same in the data; ours is delegated, and the algorithm executes, so refusing and not complying cannot be confused. Their outcome was a single accept-or-reject decision, whereas our panel separates first use, the share of wealth delegated, continuation, exit, and re-entry, so there is a commitment margin to measure at all. And they observe need through survey and portfolio measures taken at one date, whereas we build each investor's mistakes from her own trades in the months before she adopts, which is what lets us ask which mistakes predict which margin. None of the three is a criticism of their design, which was built to answer their question and does so cleanly.

We also follow adopters month by month after they join, and what happens next shows that the demand for advice is not settled at adoption. Adoption studies treat use as an absorbing state, so a product that retains assets counts as one that has worked. What the tenure evidence shows instead is a partly reversible adjustment: the tenure slope of the delegated share is near zero for the least experienced adopters and reaches -0.051 one standard deviation above mean pre-adoption experience, while the self-directed portfolio rebuilds toward its pre-adoption concentration over roughly two years. Read next to selective awareness, this makes a corrective product's success hard to distinguish from its own obsolescence, since exposure to the algorithm is followed by declines in exactly the mistakes it corrects, and delegated shares then fall fastest among the adopters best equipped to run the portfolio themselves. Both patterns connect the delegation decision to learning by trading in household finance (Linnainmaa, 2011; Nicolosi et al., 2009; Seru et al., 2010), and both imply that assets retained is the wrong summary statistic for whether such a product has worked.

3. A Framework for the Two Margins of Delegation

This section sets out a deliberately minimal framework. Its purpose is not to model why any particular investor values delegation. It is to establish two baseline results that discipline the empirical work. First, the same investor characteristic can move adoption and the delegated share in opposite directions. Second, characteristics enter first adoption, the conditional share, continuation, and the path of the share over tenure through different combinations of primitives, so no single use indicator summarizes the demand for delegation. The framework treats the loadings as reduced-form objects. The mechanism sections interpret them; the model only shows what patterns are possible and which empirical objects identify them.

3.1. Setup

Investor i has fund wealth normalized to one and characteristics X_i, the demographics, preferences, skills, and behavioral measures of the empirical sections. She may place a share $s \in [0, 1]$ of her fund wealth under delegated management. The per-period payoff from delegating share s is quadratic,

$$v_i(s) \;=\; a_i\, s \;-\; \tfrac{b}{2}\, s^2, \qquad b > 0, \tag{1}$$

where a_i is the perceived net marginal benefit of the first unit of delegation. It bundles everything the investor weighs at the margin, the expected improvement of the delegated sleeve over her own management, net of fees and of any cost she attaches to ceding fund selection. We do not unbundle it here. The perceived benefit has an observed and an unobserved component,

$$a_i \;=\; \alpha' X_i \;+\; \theta_i, \tag{2}$$

where θ_i is a latent disposition toward delegated management that the covariates do not capture.

Using the product requires a one-time entry cost,

$$k_i \;=\; \kappa' X_i \;+\; u_i, \tag{3}$$

paid at first use and expressed in per-period-equivalent units, so that flows are compared with flows.[5] The entry cost collects discovery, evaluation, and setup, the frictions that stand between a household and her first delegation decision. The central premise of the framework is that nothing ties α to κ. A characteristic can make the product easier to find and evaluate while lowering its perceived marginal benefit, or harder to find while raising it. The data, not the theory, determine the loadings.

Because the investor can always hold $s = 0$, a nonpositive perceived benefit makes non-use dominant, and entry is relevant only where $a_i > 0$. Write the per-period value of access as $V(a_i) = \max_{s\in[0,1]} v_i(s)$, which equals $a_i^2/2b$ on the range $0 < a_i \le b$ that we maintain throughout, so that shares stay interior. Conditional on entry, the investor chooses

$$s_i^* \;=\; \frac{a_i}{b}, \tag{4}$$

and she enters whenever the value of access covers the entry cost,

$$V(a_i) \;=\; \frac{a_i^2}{2b} \;\ge\; k_i. \tag{5}$$

Restricting attention to $a_i > 0$ is not innocuous bookkeeping. It is what makes the entry index monotone in the latent disposition, and the selection result below relies on that monotonicity.

[5] A forward-looking entrant would compare the discounted stream of access values, net of maintenance costs, with the one-time outlay, and with drifting benefits an option value of waiting appears. Neither refinement changes which loadings enter which margin, and that separation is all the framework is used for.

3.2. Opposite Signs Across the Two Margins

Differentiate the two margins with respect to a single characteristic x with loadings (α_x, κ_x). The conditional share moves with the benefit loading alone,

$$\frac{\partial s_i^*}{\partial x} \;=\; \frac{\alpha_x}{b}, \tag{6}$$

while the entry index in (5) moves with

$$\frac{\partial}{\partial x}\left(\frac{a_i^2}{2b} - k_i\right) \;=\; s_i^*\,\alpha_x \;-\; \kappa_x. \tag{7}$$

The benefit loading enters the entry margin scaled by the share itself. The economics of the scaling is simple. A marginal improvement in the perceived benefit of delegation is worth little at entry when the investor would delegate only a small fraction of her wealth. The entry cost, by contrast, is paid in full regardless of the share. The empirical counterpart of s_i^* is the conditional mean delegated share of 0.151, so the participation gradient is dominated by the entry-cost loading unless the benefit loading is roughly seven times larger. Opposite signs across the two margins are therefore easy to generate at realistic shares, which is the first baseline result.

Prediction 1. *The conditional-share gradient of a characteristic x has the sign of α_x, and the adoption gradient has the sign of $s^*\alpha_x - \kappa_x$. Whenever $\kappa_x > s^*\alpha_x$ with $\alpha_x > 0$, the characteristic lowers adoption and raises the conditional share. Whenever $\kappa_x < s^*\alpha_x$ with $\alpha_x < 0$, it raises adoption and lowers the conditional share. Because s^* is small, modest entry-cost loadings suffice for either pair of opposite signs.*

The latent disposition θ_i generates the selection structure the estimator is built for. On $a_i > 0$, a higher θ_i raises both the probability of entry through (5) and the share chosen after entry through (4). Participants are therefore selected on the same unobservable that drives conditional delegation, which is the sample-selection problem of Heckman (1979), and the model signs it.

Remark 1. A single latent disposition that raises the perceived benefit of delegation implies that the marginal adopter on unobservables delegates more than the average adopter, so the coefficient on the selection correction in the share equation is positive.

The remark concerns selection into the pool of active users, the margin the pooled Heckman corrects. Selection into first use prices the disposition as it stands at entry, and the dynamic extension below explains why the entry-stage and pooled corrections need not coincide.

3.3. Margins Over Time

The second baseline result is that characteristics need not affect only the decision to try the product. Let the perceived benefit drift with exposure,

$$a_{i\tau} \;=\; a_i \;+\; (\psi' X_i + \eta_i)\,\tau, \tag{8}$$

where τ is tenure, months of active use, ψ is a third loading vector, again unrestricted, and η_i is a mean-zero latent component of the drift that the covariates do not capture. The drift stands in for whatever exposure does to the perceived benefit, updating, habituation, or taste change, without taking a position on which. The share of an active user follows $s^*_{i\tau} = a_{i\tau}/b$.

Keeping the account active costs $m \geq 0$ per period, so an active user continues while $V(a_{i\tau}) \geq m$, that is while $a_{i\tau} \geq \sqrt{2bm}$, and exits when the benefit falls through that threshold. We take m small relative to entry costs, so entrants begin above it and only a downward-drifting benefit triggers exit. The size of m matters for what the model can match. When the maintenance cost is near zero, as it is for an account that costs nothing to leave open, the exit threshold sits near zero and continuation decouples from the share.

The share can then drift down for many months while exit stays rare. That is the configuration the data display, with continuation at 97.35 percent per month alongside delegated shares that decline with tenure for the most experienced adopters, and it is what lets the reclaiming-control gradient of Section 6.3 coexist with near-universal persistence. A model in which the same index governed share and survival with a high threshold could not hold both facts at once.

Each empirical margin now isolates a different combination of the loadings. First adoption depends on (α_x, κ_x) through (7) evaluated at zero tenure. The share at entry depends on α_x alone. The path of the share over tenure depends on ψ_x, so the tenure slope of delegation among active users, and its interaction with characteristics, identifies the drift; with exit rare, survivorship distorts that slope only mildly. Continuation and exit depend on ψ_x and m, since only a downward-drifting benefit crosses the maintenance threshold. Nothing restricts the loadings to share signs.

Prediction 2. *First adoption, the share at entry, the tenure slope of the share, and continuation depend on* $(\kappa_x, \alpha_x, \psi_x)$ *in different combinations. Absent proportionality restrictions on the loadings, the same characteristic can deter first adoption, raise the share at entry, and predict a declining share, with exit to follow where the maintenance threshold binds.*

The latent drift does for the dynamic margins what the latent disposition does for the static ones, and it separates the two selection corrections the paper estimates. At first use the entry index prices θ_i alone, so the model implies weakly positive selection at that stage. Among active users the share also embeds the accumulated $\eta_i\tau$, so the pooled correction prices a larger object and can exceed the entry-stage coefficient. The framework accordingly reads the small and insignificant first-adoption Mills estimate as a noisily measured, weakly positive object. A genuinely negative entry-stage correction, the cautious-experimenter reading the lifecycle section entertains, would require the disposition priced at entry to differ from the one priced in use, which the latent drift permits and the static disposition alone does not.

A composition result follows for pooled participation, the object a single use indicator measures. The stock of active users at any date mixes the entry flow with accumulated survival, and two routes produce pooled coefficients that differ in sign from flow coefficients. The first is differential survival, where a characteristic with $\kappa_x > 0$ and $\psi_x > 0$ is underrepresented among first adopters yet overrepresented among long-tenure survivors. The second requires no upward drift at all.

The characteristic itself may accumulate with use, as cumulative fund experience does mechanically, so that active months carry higher values of x than the entry months of the same investors, and

the pooled cross-section records the accumulated values. Either route lets the pooled participation coefficient carry the opposite sign from the first-adoption coefficient without inconsistency, which is the logic behind reading pooled and flow estimates side by side in Section 5.3.

3.4. What the Framework Delivers and What It Does Not

The framework delivers the objects the tables estimate. The entry condition (5) is the selection equation, the share rule (4) is the outcome equation, the latent disposition is the unobservable the inverse Mills ratio prices, the excluded platform-tenure variable is a shifter of k_i that does not enter a_i, and the drift is what the tenure and trajectory designs trace. It also delivers the two baseline results. Opposite signs across the margins are not an anomaly but the generic outcome when benefit and cost loadings differ, and characteristics shape dynamic use, not only trial, whenever the drift loading is nonzero.

Two boundaries of the framework are worth stating. Because fund wealth is normalized to one, the model speaks to delegated shares and is silent on asset levels; wealth heterogeneity re-enters there, which is why the share-based and level-based selection corrections in the robustness section need not agree. And the framework deliberately does not identify why any loading takes its sign. A negative entry gradient together with a positive share gradient is consistent with more than one story about what the characteristic does to perceived benefits and entry costs, and the drift is compatible with learning, salience, and taste change alike. This is why the framework cannot substitute for the mechanism evidence of Section 6, which brings a separate design to bear on each candidate interpretation. What the framework supplies is the mapping from loadings to the two estimating equations, and the result that opposite signs require no special assumption to arise.

4. Setting, Data, and Empirical Design

4.1. Institutional Setting

BangNiTou is a fully delegated mutual fund robo-advisor founded jointly by Ant Group and Vanguard and embedded within Alipay, China's dominant mobile payments and wealth-management platform. The product launched in April 2020 as part of China's third-party fund-advisory pilot and passed one million users within a year (Bloomberg News, 2021). The platform is itself the largest third-party mutual fund distributor in China, and the economics of fund distribution through such platforms differ from those of bank and brokerage channels (Hong et al., 2025). After collecting suitability information, it constructs and periodically rebalances a diversified mutual fund portfolio on the investor's behalf. Annual fees range from 0.20 to 0.50 percent of delegated assets, placing the service in the same broad pricing neighborhood as Betterment, Wealthfront, and other retail robo-advisors.[6] What makes the setting useful is that the same Alipay account contains three different ways of investing in mutual funds. The first is direct investing, which we label *DIRECT*: the investor picks the funds and trades them herself. The second is an automatic investment plan, which we label *AUTO*: the investor still picks the funds, but the platform buys them on a fixed schedule; selling and rebalancing remain

[6]The advisory fee is charged on delegated assets and does not depend on which funds the algorithm selects. Self-directed and automatic-plan purchases carry no advisory fee; standard fund subscription and management fees apply to fund purchases in all three modes.

hers. The third is BangNiTou itself, *RA*, where the investor hands both the fund choice and the trading to the algorithm, which selects, buys, sells, and rebalances the portfolio for her.

These three modes differ in how much of the decision the investor gives up, and it helps to think of them as rungs on a ladder. On the first rung, direct investing, she delegates nothing. On the second rung, the automatic plan, she delegates only the timing of her purchases; she still decides what to own and when to sell. On the third rung, robo-advising, she delegates the choice of what to own, and the trading that implements it, as well. The step that matters for this paper is the one from the second rung to the third. Both rungs automate the mechanics of purchase execution, so only that final step transfers the decision about which funds to hold and trade. Comparing the two therefore narrows what separates them to the investor's willingness to cede the decision itself, as opposed to her willingness to automate the mechanical work of trading. It does not reduce the difference to that single dimension, and Section 6.2 states what else the two products do not share. This comparison is the basis for the within-investor-month design in Section 6.2.

Delegation in this market is layered, and the layering sharpens the object of study. Every investor in our sample already delegates security selection to fund managers, since all of them invest through mutual funds. Comfort with delegated management as such is therefore held fixed across the sample, and the marginal decision rights, the ones the three products divide, are fund selection and the trading of the selected funds. Throughout the paper we write fund selection as shorthand for this bundle: the choice of which funds to hold together with the trading that implements it, the two rights RA takes and AUTO does not. When we speak of self-directed investing below, we mean self-directed fund picking, not stock picking, and when we speak of taking advice up, we mean transferring those marginal rights.

Because AUTO and RA sit on the same platform, they share the same account interface, the same liquidity environment, and the same execution infrastructure. The economic difference we can isolate between them is who chooses the funds, though they may also differ in branding and framing, which Section 6.2 returns to. That is what lets us study the demand for delegating a decision in isolation, without confounding it with a general taste for automation. The two products are used at comparable rates and overlap substantially: across the 4.9 million investor-months, 13.95 percent record both RA and an automatic plan, 18.12 percent record RA only, 18.99 percent record an automatic plan only, and 48.94 percent record neither.

This study was remotely conducted on the Ant Open Research Laboratory [7] in an Ant Group Environment. All data were sampled and desensitized, and were analyzed on the Ant Open Research Laboratory. The laboratory is a sandbox environment where the authors can only remotely conduct empirical analysis and individual observations are invisible.

4.2. Sample Construction and Estimand

The sample contains 100,000 investors observed monthly for 49 months from September 2020 through September 2024, which yields 4.9 million investor-months. To enter the sample, an investor had to be registered on Alipay before the sample start, fall in the eligible age range at sampling, and have at least one consumption record, mutual fund holding, purchase, and redemption during the observation window. The activity requirements are not incidental. The behavioral measures are built from each

[7] www.deor.org.cn/labstore/laboratory

investor's own fund trades, so an investor with no trades would have no measurable biases, and the questions of the paper are questions about active choices. The resulting sample is young relative to typical Western brokerage samples. Mean age is 32.5, median age is 30, and the range runs from 19 to 64. Female investors make up 42.5 percent of the sample. Mean on-platform total financial assets are CNY 95,245, with a median of CNY 30,220. These moments bound the interpretation, since the paper studies active mobile-platform fund investors, not the full household population.

The sample is stratified on ever adoption, meaning that an investor's chance of being drawn depends on whether she ever used RA during the window. Of the investors in the summary statistics, 70,034 use RA at least once and 29,966 hold mutual funds but never use RA. Conditional on ever adopting, investors hold RA for 22.4 months on average. This oversampling of adopters is deliberate. In the underlying population, first adoption is a rare monthly event, and the intensive-margin outcome exists only in months when the investor holds positive RA assets. A uniform draw of the same size would therefore contain too few first-adoption and RA-active observations to estimate the lifecycle and conditional-share relations with any precision.

Stratifying this way fixes what the sample can and cannot tell us. It cannot tell us the platform-level prevalence of RA use, since the stratification mechanically inflates the adopter share. What it can tell us is how investor characteristics relate to each margin of the participation decision, among the active mutual fund investors on the platform. That cross-sectional slope is the estimand. The distinction is standard in choice-based sampling for rare outcomes (Manski & Lerman, 1977), where the intercept and the unconditional prevalence are not recoverable but the slope of characteristics within the sampled strata remains informative. The behavioral questions in this paper are questions about slopes, so the design serves them.

A last design choice concerns weighting. We report unweighted estimates throughout the main text and reserve the weighted estimates for robustness. Population weights would re-scale the ever-adopter and never-adopter strata back toward their platform shares, which is the right correction when the target is a population prevalence. Our target is a slope rather than a prevalence, and we define it on the sampled population: the relation between an investor characteristic and each margin among the 70,034 ever-adopters and 29,966 never-adopters we observe. We do not claim that the unweighted slope recovers the platform-population slope without further assumptions. Stratification here is on ever adoption over the full window rather than on the monthly outcome, the selection stage is a probit rather than a logit, and ever-adopters contribute pre-adoption, active, and post-exit months, so the case-control results that leave logit slopes unaffected do not transfer. What the weighted re-estimation establishes is that the behavioral sign pairs are the same under either estimand. The weighted re-estimation in Online Appendix Table A.1 preserves the cross-margin sign pairs that carry the paper's argument, so none of the behavioral findings turns on the choice; the selection term is the one coefficient that does not survive reweighting, and Section 7.5 says what that does and does not imply. We report the unweighted estimates as the main results because they are the more precisely estimated version of the same relations. We do not use either set to infer unstratified platform penetration.

Two further features of the design follow from the stratification. First, investors are sampled on ever adoption, not on adoption in a given month, so each ever-adopter contributes her entire monthly history, inactive months included. The monthly Fama–MacBeth estimates then draw on

cross-sectional variation within each month while preserving each investor's lifecycle path. Second, never-adopters enter the selection equation but not the outcome equation, since the conditional RA share is undefined when RA assets are zero. For this reason we report extensive- and intensive-margin moments separately, and we use the conditional RA-share mean as the denominator when we translate outcome-stage coefficients into economic magnitudes.

4.3. Outcomes and Covariates

Three margins organize the analysis throughout. The extensive margin is the binary decision of whether to adopt robo-advising at all. The intensive margin is the share of fund wealth an investor delegates once she has adopted. The lifecycle margins are first adoption, continuation, exit, and re-entry. In the terms of the introduction, the extensive margin is entry, and the intensive and lifecycle margins are the two views of commitment, how much the investor delegates and how long she leaves it delegated. The main extensive-margin outcome is an indicator for holding positive RA assets in month t. The main intensive-margin outcome is

$$s_{it} = \frac{\text{RA Value}_{it}}{\text{RA Value}_{it} + \text{Mutual Fund Value}_{it}}, \qquad d_{it} = 1. \tag{9}$$

The denominator adds the robo sleeve to the self-directed sleeve, the part of the fund portfolio the investor still manages herself. The share therefore measures how much of her fund wealth she has handed to the algorithm. Across investor-months where the denominator is defined, the unconditional mean of this share is 0.053. Among the 1,571,475 RA-active investor-months, the conditional mean is 0.151. The conditional mean is the figure we use when interpreting outcome-stage coefficients. Appendix specifications report two alternative intensive outcomes, log RA assets and RA assets over total financial assets.

Three lifecycle outcomes separate states that a single binary RA-use measure would pool together. First-adoption events equal 1.43 percent of all panel investor-months. That is an incidence per observed month rather than a transition rate, since its denominator includes the post-adoption months of investors who have already entered; the hazard among not-yet-adopters is higher, and Section 5.3 reports the 3.2 percent figure that conditions on at-risk months. Continuation, the rate of still holding RA given that the investor held it the prior month, is 97.35 percent across the 1,532,289 at-risk months. Re-entry among lapsed adopters is 1.30 percent across their 752,774 at-risk months. These three rates describe economically different situations. First entry is rare and reflects discovery and initial evaluation. Continuation dominates the RA-active panel and reflects persistence in a product the investor already knows. Re-entry is a return after a match that earlier ended or paused. A single use dummy would blur all three, and the lifecycle analysis therefore treats them separately.

Covariates fall into demographics, platform assets, risk preference, prior robo-product exposure, platform tenure, fund experience, selection ability, timing ability, platform engagement, and ten behavioral measures. The demographic and account controls are log total financial assets, age, gender, a survey-based six-level risk preference, cash and money-market balances, and an indicator for using other robo products on the platform. The cash and money-market balances are Yu'e Bao and similar platform cash-management holdings. We read them as a joint proxy for liquidity preference,

precautionary saving, and payment-account intensity, not as risky wealth.

Fund experience is the cumulative number of months in which the investor held at least one mutual fund; it averages 23.2 months in the panel. It is a stock of experience, not a measure of recent returns. The distinction carries into our interpretation later on, which turns on accumulated familiarity with selecting and monitoring funds; realized returns play no part in it. Selection ability and timing ability split an investor's realized fund returns into two parts: what she earned by holding the right funds, and what she earned by trading them at the right moments. Selection ability is the value-weighted excess return of the investor's funds relative to their Morningstar categories; in plain terms, it asks whether she holds funds that beat their category peers. Timing ability is the profit and loss from within-month trading, net of the passive return on her beginning-of-month holdings, so it asks whether she buys and sells at the right moments. The two can diverge, and a single performance number would conflate them, and we therefore separate them following Grinblatt and Keloharju (2009).

The ten behavioral biases are central to the analysis and come from separate strands of the behavioral-finance literature, so we first say in plain terms what each is meant to capture and turn to measurement afterward. Underdiversification is the tendency to hold too few or too similar assets, so that the portfolio carries avoidable risk that diversification would remove at no cost in expected return. Total turnover is the rate at which the investor churns her portfolio. The revealed-measure literature on overconfidence works with several trading-based proxies, of which total turnover is the most common (Barber & Odean, 2000, 2001), since investors who overestimate the precision of their information trade more than is warranted. Speculative turnover is a second such proxy, and a distinct measure rather than a component of total turnover's interpretation. It counts the round trip in which a position is sold soon after a gain and the same or a similar fund is quickly rebought, so it captures deliberate position-level speculation, whereas total turnover captures how much the investor trades overall. The two are constructed to be separable, since an investor can churn heavily without ever completing a round trip, and they enter the specifications together.

The remaining seven capture patterns delegation has no obvious counterpart for. Limited attention is the tendency to trade whatever funds are salient, such as those with extreme recent returns or fresh launches, instead of those a deliberate search would surface. The disposition effect is the tendency to sell winners too early and hold losers too long. Extrapolation is return-chasing, buying funds because their recent returns were high. Gambling preference is a taste for lottery-like payoffs, cheap funds that offer a small chance of a large gain. Anchoring is the tendency to let salient reference prices, such as a fund's twelve-month high or low, drive the decision to trade. Confirmation bias is the tendency to hold on to positions whose past performance has validated the investor's earlier judgment. Herding is the tendency to buy and sell what the crowd is buying and selling. Each is costly or distortionary in the sense the cited literature establishes, and each is built to separate the behavioral pattern from its rational counterpart.

Each measure is constructed to capture its own mechanism, not a single underlying trait, and Online Appendix A.4 gives the construction of each. Underdiversification is portfolio HHI, the sum of squared portfolio weights, which averages 0.423 in the panel; the mean investor-month effective number of funds, computed as $1/\mathrm{HHI}_{it}$ and then averaged, is 4.9; it connects to evidence that retail investors hold concentrated portfolios and often miss the diversification cost (Blume & Friend, 1975;

Goetzmann & Kumar, 2008; Huberman, 2001). Total turnover is the most common of the revealed overconfidence proxies, linked to evidence that high-turnover retail investors underperform on average (Barber & Odean, 2000, 2001). Attention bias captures trading in salient funds, in the spirit of Barber and Odean (2008); the disposition effect follows the PGR-minus-PLR logic of Odean (1998); extrapolation captures lagged return-chasing (Barberis et al., 2015; Greenwood & Shleifer, 2014); gambling bias identifies trading in lottery-like funds (Kumar, 2009); anchoring captures trading near prior price extremes; confirmation bias measures retention of lagged winner positions; and herding measures trading in funds with high lagged aggregate flow, in the spirit of Lakonishok et al. (1992). Three construction principles keep the behavioral patterns distinct from their rational counterparts. Bias measures are lagged when we model the month-t decision, directional components are separated where possible, and aggregate-flow variables exclude the investor's own flow. These choices are what let Section 6.1 tell correctable active-opinion biases, attachment biases, and cognitive-constraint biases apart; a noisier omnibus behavioral index would not support that classification.

Continuous regressors are winsorized at the first and ninety-ninth percentiles and standardized within month. Winsorizing limits the influence of extreme platform records, and within-month standardization puts every coefficient in per-standard-deviation units that are comparable across months, which is the unit the Fama–MacBeth average requires. Two conventions follow and hold throughout the paper. A one-standard-deviation change in a regressor means one within-month standard deviation, which need not equal the pooled standard deviation reported in Table 1, since the pooled figure also carries variation across months. And outcomes stay in their natural units unless a table's notes say otherwise, so a coefficient on the delegated share is a change in that share, whereas the dose-response outcomes of Section 7.2 are z-scored within month and are read in standard deviations of the outcome. Binary variables are left unstandardized. All right-hand-side variables except age, gender, and risk preference are lagged one month. Bias variables are forward filled across non-trading months once the investor has made her first observed trade, on the premise that a behavioral tendency persists until new trading evidence updates them. Returns are never forward filled.

4.4. Heckman Specification

The baseline design estimates a Heckman two-stage model separately in each monthly cross-section. The first stage is a probit for RA participation:

$$\Pr(d_{it} = 1 \mid Z_{it-1}) = \Phi(Z'_{it-1}\gamma). \tag{10}$$

The second stage is an OLS outcome equation for the RA share among RA-active investor-months:

$$s_{it} = X'_{it-1}\beta + \delta\widehat{\lambda}_{it} + \varepsilon_{it}, \qquad d_{it} = 1, \tag{11}$$

where $\widehat{\lambda}_{it}$ is the inverse Mills ratio. The inverse Mills ratio is a generated regressor built from the first-stage probit, and it captures the part of the adoption decision the observed covariates leave unexplained. We estimate the cross-sectional relationship within each calendar month and then average the monthly slopes, which gives Fama–MacBeth coefficients (Fama & MacBeth, 1973). Standard errors are Newey–West with lag three, which allows the monthly coefficient estimates to be serially

correlated within a quarter. They do not carry the first-stage estimation error, since the second stage treats the inverse Mills ratio as observed. With monthly cross-sections of this size the first-stage index is estimated precisely, so that omission is second order relative to the cross-sectional variation the coefficients use; a Murphy–Topel adjustment or a bootstrap would be the formal correction. Province dummies enter both stages to absorb regional differences in financial development, platform penetration, and product diffusion. The baseline, skills, and biases specifications retain 48, 47, and 36 monthly cross-sections; the bias panel is shorter because its behavioral measures require rolling windows that the earliest months cannot fill.

The Heckman correction is part of the economic design here, not just a statistical fix. The two margins are linked, because the delegated share is observed only after the investor has chosen to adopt. The same unobserved disposition, a latent trust in algorithmic management or comfort with ceding control, plausibly raises both the probability of entry and the share delegated once inside. If it does, then the investors we observe in the RA-active subsample are selected on the very dimension that also drives the outcome. An OLS regression run on RA-active months alone would then confound the observable slopes with that selection. Modeling adoption and conditional delegation jointly is the approach this structure calls for, since the second margin is realized only after the first. The inverse Mills coefficient is, in turn, something to read rather than something to sweep up. In the two-step estimator it equals $\delta = \rho\sigma_\varepsilon$, so its sign is the sign of the correlation between the unobservables of the two equations, and it says whether the investors who select into RA on unobserved dimensions go on to delegate more or less once they enter. Treating that coefficient as something to interpret, rather than as a nuisance to be swept up, is standard in the applied selection literature. That literature also sets the conditions under which the interpretation is allowed. There must be a variable that shifts entry and is left out of the share equation, an argument for what the unobserved trait is, and evidence that the term is not standing in for an observable the model omitted (Heckman, 1979; Li & Prabhala, 2007; Vella, 1998). Section 5.2 meets those conditions in turn rather than asserting them.

The two-step estimator does not constrain the implied $\hat{\rho} = \hat{\lambda}/\hat{\sigma}$ to the admissible range $[-1, 1]$, and in the biases specification our average $\hat{\rho}$ is 1.381. We therefore report $\hat{\lambda}$, which is what the regression estimates, rather than $\hat{\rho}$, which the parameterization would have to deliver for a literal error correlation to be recovered. The practical reading is that the correction is doing too much work: the inverse Mills ratio carries more of the outcome variation than a bivariate-normal model with these regressors can attribute to a correlation alone. Signs and relative magnitudes on the observable slopes are what we interpret, and we do not read $\hat{\lambda}$ as a calibrated correlation.

Estimating month by month, rather than pooling, is substantive for a related reason. Many covariates move slowly or not at all, including age, gender, baseline wealth, risk preference, and accumulated experience. Investor fixed effects would absorb much of the cross-sectional variation that the main question depends on. Estimating monthly cross-sections keeps that variation intact, lets the coefficients differ across market environments, and produces a time-series average that summarizes the average cross-sectional relation over the 49-month window. This choice does give up the ability to difference out time-invariant investor heterogeneity. To address that cost, Online Appendix Table A.5 reports a panel fixed-effects control-function alternative, in which the three focal behavioral intensive-margin signs are unchanged while the demographic gradients are, by construction, absorbed into the

investor fixed effect. Two of the smaller bias coefficients, anchoring and gambling, do change sign under that estimator. Neither carries an argument in this paper, and we read them as the coefficients least well determined by the data rather than as findings.

Platform tenure serves as the excluded variable, and the exclusion rests on a discovery channel. Longer platform tenure raises the probability of encountering and evaluating RA, and the relevance condition is strong: months since registration enters the selection equation with marginal effects of 0.011, 0.008, and 0.014 in the baseline, skills, and biases panels, with t-statistics of 10.54, 7.80, and 10.98. Two things support the exclusion itself, and neither uses output from the corrected model. Province dummies absorb the regional differences in platform penetration that would otherwise make tenure a proxy for local diffusion. And the characteristics that carry the paper's argument are pre-adoption trading histories, measured on the investor's own fund transactions rather than on her relationship with the platform, so tenure and the regressors of interest are built from different records.

The objection we take most seriously is that tenure proxies for attachment to the platform, since provider tenure is a standard determinant of intensive-margin attachment through switching costs (Klemperer, 1995). Under that reading tenure belongs in the outcome equation and the exclusion fails. Online Appendix Table A.2 puts tenure in both equations and, separately, in neither, and Section 7.5 reads what those arms do and do not establish. As in every Heckman design the exclusion is not directly testable, and it is the main identifying restriction behind the selection correction, which is why Section 4.5 sets out what the paper's central claim does not require it for. Selection-stage coefficients in the non-interacted Heckman tables are reported as approximate marginal effects, evaluated at each panel's own RA-active rate; Online Appendix A.2 gives the scaling. Outcome-stage coefficients are reported in raw RA-share units, so a coefficient of 0.020 means a two-percentage-point increase in the conditional RA share, about 13 percent of the 0.151 conditional mean.

4.5. What the Two-Margin Result Does Not Rest On

The exclusion restriction is the design's one untestable assumption, so it is worth setting out how much of the paper would stand if a reader rejected it. The answer is most of it, for three reasons that can be checked in the appendix tables. The adoption margin does not use the exclusion restriction at all. The selection equation is a probit that includes platform tenure as a regressor like any other, and dropping it changes nothing: across the arms of Online Appendix Table A.2, concentration enters the selection equation at -0.172 with tenure included and -0.170 with tenure dropped, turnover at -0.083 against -0.084, and confirmation at $+0.080$ against $+0.081$. Half the paper's central pattern, the half that says the investors with the most correctable mistakes are the least likely to adopt, is therefore invariant to where tenure is placed.

The intensive margin is where the exclusion does work, and one design reproduces its sign pattern without using the exclusion at all. The authority ladder of Section 6.2 contains no selection correction of any kind: it stacks the three product rows within the investor-month and identifies the automation and delegation steps from within-investor-month variation, with no first stage and so no excluded variable. It delivers the same split, with concentration raising the automation step and lowering the delegation step. Two further designs keep the same exclusion but change the variation it is applied to, which is weaker than dispensing with it and still informative. The first-adoption design of Section 5.3

measures the flow into use from pre-adoption histories and recovers the negative-then-positive pair at the investor's first delegation decision, so the pattern does not depend on the stock of active users. The panel control function of Online Appendix Table A.5 replaces Fama–MacBeth with investor fixed effects and two-way clustering and returns concentration at $+0.045$, turnover at $+0.040$, and confirmation at -0.015 on the conditional share, so the pattern does not depend on cross-sectional averaging either.

Nor is the positive share coefficient a property of the denominator. The delegated share divides RA assets by RA plus mutual fund assets, so a mechanical reading would attribute the positive sign to concentrated investors holding fewer funds and therefore a smaller denominator. Two alternative outcomes rule that out. Dividing RA assets by total financial assets, a denominator the concentration measure does not enter, leaves concentration at $+0.017$ and turnover at $+0.017$ with confirmation at -0.006. Using the log level of RA assets, which has no denominator at all, leaves them at $+0.204$, $+0.069$, and -0.048 (Online Appendix Table A.4). What the exclusion restriction does carry is the magnitude of the selection term itself, and with it any quantitative statement about how much of the delegated share is attributable to unobserved selection. Section 7.5 reports the sensitivity and reads it.

5. Who Uses Robo-Advising?

Throughout this section we read each magnitude against the relevant base rate. Extensive-margin coefficients are percentage-point changes in the probability of RA use, measured against the table-specific participation mean, which runs from 32.6 percent in the baseline panel to 37.2 percent in the biases panel, and which is 32.3 percent in the return specifications of Section 6.4. Intensive-margin coefficients are percentage-point changes in the conditional RA share, measured against the 0.151 mean delegated share among RA-active investor-months. The distinction is not cosmetic, since the same coefficient carries different economic weight on a 32-to-37 percent participation base than on a 15.1 percent conditional-share base. And because the sample is stratified on ever adoption of RA, these benchmarks describe the estimation sample, not the platform-population prevalence.

5.1. Baseline Heckman Estimates

Table 2 reports three nested Heckman specifications. The baseline panel includes demographics and account variables. The skills panel adds fund experience, selection ability, and timing ability. The biases panel adds the ten behavioral measures, with total and speculative turnover entering together. The three panels retain 48, 47, and 36 monthly cross-sections. Among the demographics, wealth operates almost entirely at entry. A one-standard-deviation increase in log total financial assets raises RA participation by 14.2 percentage points in the baseline panel, 13.4 in the skills panel, and 7.9 in the biases panel. Measured against the table-specific RA-active rates, these are roughly 44, 39, and 21 percent of the relevant participation means. On the intensive margin the picture is different. The conditional-share coefficient is negative in the baseline and skills panels, at -0.079 and -0.049, which is about 52 and 32 percent of the 0.151 conditional mean. But once behavioral controls enter, it falls to -0.010, about 7 percent of the conditional mean, and loses significance. Wealth therefore has a large and robust effect on entry, while its baseline relation to delegated share is mostly absorbed

by behavioral composition. This pattern helps reconcile a disagreement in the prior literature over whether larger portfolios substitute away from robo-advising. Larger portfolios enter more readily, but their conditional commitment is not lower once behavior is held fixed.

Risk preference, by contrast, raises both margins at once, by 1.9 to 2.9 percentage points on participation, 5 to 9 percent of the participation mean, and by 0.8 to 2.0 on the conditional share, 5 to 13 percent of its mean. It is the table's benchmark for a characteristic that does not separate the two margins, which is what makes the opposite signs below attributable to those variables rather than to the two-stage machinery.

Female investors are less likely to use RA and delegate less once they do, by four to five percentage points and 1.3 to 2.0 respectively, roughly 12 to 14 percent and 9 to 13 percent of the two means, in every panel. Prior use of the platform's other robo products is the strongest observed predictor of participation, at about 22 percentage points, between 59 and 69 percent of the participation mean, with a conditional-share coefficient rising to 6.7 percentage points in the biases panel. That variable carries persistent heterogeneity in willingness to use delegated products which wealth, risk preference, and the biases do not exhaust.

Age shows why the two-margin design matters for reading demographic evidence. In the baseline specification, age has little adoption effect and a small negative conditional-share coefficient of -0.005. After behavioral controls enter, age becomes a positive adoption predictor, at 1.6 percentage points per within-month standard deviation, while its conditional-share coefficient is essentially zero and insignificant. The reason is that age bundles offsetting behavioral channels. Online Appendix Table A.17 shows that older investors trade less, hold less concentrated portfolios by HHI, and display stronger confirmation bias. Once those channels are separated out, age is best read as a conditional extensive-margin effect. The post-bias coefficient is modest next to the wealth and prior-robo coefficients, but it is not negligible, at about 4 percent of the participation mean in the biases panel. This is one explanation of why binary adoption studies reach different conclusions about age, since the sign depends on which correlated behaviors a given design happens to observe.

Fund experience gives the main skill-side cross-margin pattern in the pooled Heckman. In the skills panel, fund experience raises RA participation by 3.5 percentage points and lowers the conditional RA share by 2.8 percentage points. The adoption effect is about 10 percent of the skills-panel participation mean, and the intensity effect is about 19 percent of the conditional mean, so both are economically meaningful and not merely significant. The positive adoption sign fits experience lowering the cost of evaluating the product, since an experienced fund investor understands the asset class and can size up a fund-of-funds robo product easily. The negative intensity sign is the one that needs an account, and more than one is available at this stage. Experience may raise what an investor forgoes when she cedes fund selection, or it may proxy for behavioral correlates that carry the relation, or the pooled coefficient may reflect who survives in the active pool rather than who commits. We do not choose among them here. Section 5.4 sets the candidates out, and Sections 6.2 and 6.3 bring separate designs to bear on them. The joint specification in Online Appendix Table A.3 keeps the positive participation coefficient, at 2.4 percentage points, but attenuates the conditional-share coefficient to -0.005 and insignificance. Experience therefore has an independent extensive-margin component, while part of its intensive-margin relation runs through behavioral correlates. Table 3 adds a lifecycle qualification,

since pre-adoption fund experience lowers first RA adoption. Experience can deter the initial move into delegation even as it predicts being found in the RA-active pool once the product has been used.

The behavioral variables generate the paper's central taxonomy. Underdiversification and total turnover share what we call the awareness-correctable signature. A one-standard-deviation increase in underdiversification lowers RA participation by 6.5 percentage points and raises the conditional RA share by 2.0 percentage points. The entry effect is about 17 percent of the biases-panel participation mean, and the share effect about 13 percent of the conditional mean. The interpretation rests on the pair of signs rather than on either coefficient alone. A concentrated investor could in principle show any of four combinations, and only one is the awareness signature. Negative on both margins would mean simple distaste for the product. Positive on both would mean full recognition of the cost and proportional commitment. Positive-then-negative would be the attachment pattern we document below for confirmation. The observed negative-then-positive split is the one the awareness reading predicts, and Section 5.4 takes up the readings that could produce it too. The investors who would benefit most are slower to adopt; those who do adopt delegate heavily, and their concentration then recedes over tenure, in the direction a diversified algorithmic portfolio making the cost visible would imply. Turnover tells the same story with the weight shifted. It lowers participation by 3.0 percentage points, about 8 percent of the participation mean, and raises the conditional share by 4.5 percentage points, about 30 percent of the conditional mean. This is the same negative-then-positive signature, but tilted toward intensity, not entry, the opposite weighting to underdiversification. Overtrading deters trying the product less, plausibly because high-turnover investors are already active, but it matters more for how much they hand over, since the rule-based product substitutes most directly for the discretionary trading the measure captures. Both sign pairs are stable across every check that retains the exclusion restriction, including the joint specification, population reweighting, and a panel control-function estimator with investor fixed effects; Section 7.5 summarizes the set and the Online Appendix reports the tables. The economic logic is demand-side selection into corrective technology. Concentrated and high-turnover investors make mistakes that diversified, rule-based RA can discipline (Barber & Odean, 2000, 2001; Goetzmann & Kumar, 2008), yet the same investors may be the least likely to perceive the cost before they adopt.

Speculative turnover separates from total turnover on the intensive margin, and the separation needs its own discussion, because the two are often treated as interchangeable proxies for overconfidence and here they pull the delegated share in opposite directions. Holding total turnover fixed, speculative turnover lowers participation slightly, by 0.2 percentage points ($t = -3.37$), and lowers the conditional share by 1.9 percentage points ($t = -8.35$), about 13 percent of the conditional mean, against total turnover's positive 4.5.

What separates them is whether the robo-advisor can substitute for the behavior. RA trades rarely, since it holds a diversified portfolio and rebalances on a schedule, so an investor who delegates a large share simply stops making most of the trades she used to make. Total turnover is therefore something delegation can displace. Much of that churn is undirected activity with no particular position behind it, the trading that Barber and Odean (2000) show costs retail investors their returns, and handing it to a rule-based portfolio removes most of it. Heavy traders hand over more for exactly that reason, and their turnover then falls further with tenure, which is the trajectory evidence of Section 6.1.

Delegation cannot take over speculative turnover in the same way. It counts a directed act, the sale of a recent winner followed by a quick repurchase, and the investor makes that trade because she believes she knows when to be out and when to be back in. This is overconfidence in the sense the literature means it, a belief in the precision of one's own signal (Gervais & Odean, 2001; Statman et al., 2006), and it is expressed at the level of a position rather than in the volume of trading. A rule-based product has no way to act on such a belief, so the wealth an investor wants managed that way stays outside RA and her delegated share is lower. Delegation absorbs the overtrading without touching the confidence that generates the directed trades. Table 4, which Section 6.2 introduces, shows the same thing from another angle: speculative turnover moves investors away from the automatic plan and away from the robo-advisor alike, at -0.031 and -0.013, so they decline to hand the decision to either product. The distinction holds up under reweighting and the other checks the appendix collects alongside the behavioral measures.

The split inside turnover is a small instance of what separates the rows of this table from one another. Delegation substitutes for behavior that consists of doing what the algorithm also does, and leaves untouched the beliefs and constraints an investor carries into the product. Trading volume and portfolio construction belong to the first group, which is why underdiversification and total turnover raise the delegated share, and why exposure should reduce them if the account is right. Position-level conviction belongs to the second, and so do limited attention, a taste for lottery-like payoffs, and herding, each of which lowers both margins because delegating changes neither what an investor attends to nor what she wants to own. Attachment to validated positions sits between the two, easy to try around and hard to hand over. Section 6.1 develops these three signatures and tests them on post-adoption trajectories; what the turnover split adds here is that the division can run inside a single bias measure, so that one label holds a behavior RA replaces and one it cannot.

Confirmation bias runs the other way. It raises participation by 3.0 percentage points and lowers the conditional RA share by 2.2 percentage points, about 8 percent of the participation mean and 15 percent of the conditional mean. The pair survives the joint specification of the appendix. The pattern fits an attachment interpretation, and its positive-then-negative ordering is the exact opposite of the awareness biases, which is what makes it informative. An investor whose prior choices have recently been validated can try a separate RA sleeve without unwinding her own positions, so entry is easy. But a large RA share would require capital to leave the positions or narratives that confirmation supports, so commitment is held back. That two biases sort investors in opposite directions across the margins is also why a single behavioral index would have averaged them toward zero and missed both. Attachment to one's own positions is what the disposition-effect literature documents (Odean, 1998; Shefrin & Statman, 1985), including among mutual fund investors (Bailey et al., 2011), and delegation is known to change it: the disposition effect holds for self-chosen assets but takes the opposite sign for delegated ones, where the manager can be blamed (Chang et al., 2016), and the attachment reading extends that logic from realizing losses to committing capital.

Attention, gambling, and herding biases reduce both margins, which is what the framework of Section 3 implies for a characteristic that loads on the entry cost without raising the perceived benefit of delegating. Their magnitudes are smaller than the main active-opinion effects, between 0.6 and 0.9 percentage points on entry and between 0.3 and 0.5 on the conditional share, but the same-signed pat-

tern is itself informative. A bias that lowers both margins is one whose underlying constraint exposure to the product does not relax. The underlying constraint travels with the investor into the robo sleeve instead of being corrected by it, so there is no margin on which awareness can show up. The split between this same-signed group and the opposite-sign awareness group is the paper's internal falsification test, since a model in which all biases responded uniformly to RA would predict one common pattern. The result fits cognitive constraints or preference-like behaviors that RA does not necessarily relax (Barber & Odean, 2008; Kumar, 2009; Lakonishok et al., 1992). The remaining three biases fit neither pattern and we do not force them into either; the table reports anchoring, extrapolation, and the disposition effect, and the argument leans on none of them. What the panel as a whole shows is that behavioral biases are heterogeneous in their demand signatures, dividing into correctable active-opinion biases, attachment-type biases, and cognitive constraints that do not respond to delegation in the same way.

This baseline evidence speaks directly to the conflicting RA-adoption literature. Studies that describe adopters as younger, male, and smaller-portfolio holders are estimating a mixture of entry and commitment, in which demographics and wealth absorb unobserved behavioral composition. Studies that find weak demographic effects are closer to specifications in which those behavioral channels are observed or less dominant. The present table shows why both sets of findings can hold within a single platform. Wealth is strongly positive on entry but not on conditional delegation once behavior is controlled. Age operates conditionally at entry and hardly at all on intensity. And behavioral biases carry heterogeneous signs across the two margins. The reconciliation is therefore not that one prior result is wrong, but that a binary adoption outcome is too coarse for a delegated product.

5.2. Selection on Unobservables

Selection on unobservables is positive in all main panels. The Mills coefficient is 0.158 in the baseline panel, 0.203 with skills, and 0.212 with biases, and the joint specification gives 0.277. These are large coefficients on a share outcome whose conditional mean is 0.151. The coefficient should not be read one-for-one as a percentage-point effect, since the inverse Mills ratio is itself a generated nonlinear index. But its persistence across specifications shows that the unobserved adoption-intensity channel is not small relative to the observed slopes in the same table.

What matters economically is the sign, not the size. A positive coefficient means that the marginal adopter, the investor drawn in by something the covariates do not capture, delegates more than the average adopter, not less. Two readings can be ruled out on that basis. A cautious-experimenter reading would predict a negative coefficient, since an investor testing the product cautiously would delegate less than average. A pure adverse-selection reading would also point negative. The positive sign fits neither. What fits is one trait, unevenly distributed across investors, that pushes the same way twice: it makes them more likely to adopt and, once they have, more willing to hand over a large share. Comfort with letting an algorithm manage money is the natural candidate.

The coefficient's stability across specifications says something more. Adding fund experience, selection ability, timing ability, and the behavioral measures does not absorb the selection channel; if anything the coefficient rises, from 0.158 to 0.212. If the observed skills and biases were sufficient statistics for the unobserved willingness to delegate, the Mills coefficient would shrink sharply once

they entered. It does not. The residual channel is therefore not simply underdiversification, turnover, confirmation, or measured performance skill in unobserved form.

Naming that disposition is where the applied selection literature counsels restraint, and we follow it. The inverse Mills coefficient identifies a covariance between the unobservables of two equations, not a variable, and the conditions under which it can be given economic content are an excluded shifter that moves entry, a stated argument about what the unobservable is, and evidence that it is not an omitted observable in disguise (Li & Prabhala, 2007; Vella, 1998). The stability just described meets the third condition. Two further patterns in the appendix bound what the disposition can be.

It is not a wealth or scale proxy. The same selection term is positive on the share outcomes and negative on the level outcome, at -1.411 on the log of RA assets against $+0.212$ on the delegated share (Online Appendix Table A.4). An unobservable that made adopters larger would raise both. One that makes an investor willing to place a larger fraction of her risky sleeve under the algorithm raises the share while leaving the level alone or lowering it.

That is as far as the design reaches, and the trust-in-finance literature supplies the name rather than the identification. Trust predicts participation beyond wealth, education, and measured sophistication (Guiso et al., 2008), and delegated management can offer reassurance as much as alpha (Gennaioli et al., 2015); trust has also been proposed as a determinant of robo-advice demand specifically (Belanche et al., 2019; Scherer & Lehner, 2023). That is the trait those theories describe: something an investor carries with her that changes how she splits her money between herself and a manager, without changing how much money she has. The platform setting makes a durable version of it plausible, since a user who has trusted Alipay as a financial interface for years may be readier both to try RA and to place a larger share inside it. We use the word in that sense.

A second reading predicts the same coefficient and has to be separated from it. Selective awareness also says that the investors who cross the entry hurdle commit deeply, so an unmeasured component of awareness would show up here as a positive Mills coefficient and could be mistaken for trust. The two are nonetheless different objects, and the difference is testable in principle. Awareness is a belief about the investor's own portfolio, so it should be moved by exposure to the algorithmic one, and Section 6.1 shows that the mistakes it concerns do recede after adoption. Trust is a belief about the product, and nothing in our design makes it a function of what the investor sees in her own holdings. What we can say is that the residual is not the measured part of awareness, since underdiversification and turnover are in the specification and do not absorb it. What we cannot say is how much of the residual is unmeasured awareness rather than trust, and we do not claim a decomposition. The design does not separate latent trust, platform attachment, and a general affinity for delegated services either. This residual channel is part of the contribution. The headline coefficients show that observed characteristics can move adoption and delegated share in different directions. The Mills coefficient shows that observably similar adopters also differ on an unobserved dimension that raises conditional delegation. A binary adoption study would fold both channels into one reduced-form use coefficient. The two-stage design instead makes the observable cross-margin signs and the residual selection channel separately visible.

5.3. First Adoption

One reading of the baseline result cannot be ruled out with pooled data. If concentrated investors who happen to remain in the RA-active pool are the ones delegating heavily, the positive share coefficient records survivorship and not a response to the product. Testing that requires the flow into use rather than the stock of users. The pooled Heckman measures a stock, the set of investors observed holding RA in a given month. A stock is not the same as the flow into it, and several of the paper's characteristics could in principle act on the two differently. A trait that makes an investor more likely to be found among active users could, at the same time, make her slower to take the first step. The pooled design cannot tell these apart, because every active month enters it regardless of when the investor first adopted. To isolate the first step we need a design built on the flow. Table 3 provides one. It reports first adoption of RA and AUTO using a pre-window construction, in which the regressors summarize the investor's pre-adoption history, not her contemporaneous active-user state.

Keeping the stock and the flow straight is what turns apparent sign differences into lifecycle heterogeneity. Pooled participation records who is observed holding RA in a given month, and that pool mixes first adopters, continuing users, re-entrants, and long-tenure survivors. First adoption records something narrower, namely who crosses the entry threshold for the first time, out of the at-risk pool of investors who have not yet adopted. When the pooled and first-adoption signs differ, then, the gap is not a contradiction. It is information about which part of the lifecycle a characteristic acts on.

Wealth raises first adoption of both products and lowers the initial conditional share. For RA, the first-stage marginal effect of pre-window log total financial assets is 0.029 and the initial RA-share coefficient is -0.041. For AUTO, the corresponding coefficients are 0.042 and -0.080. These entry effects are economically meaningful, since the RA first-adoption rate in the table is 7.7 percent and the AUTO rate is 8.4 percent. The wealth coefficient is about 38 percent of the RA first-entry rate and about 50 percent of the AUTO first-entry rate. That 7.7 percent needs explaining, since it is higher than the 1.43 percent of panel investor-months in which a first adoption occurs (Table 1). The two rates have different denominators. The pre-window design collapses each never-adopter's history into a single at-risk observation, so its denominator is the pool of pre-adoption and never-adopter records instead of all investor-months. The initial-share effects are material too, about 27 percent of the conditional RA-share mean for RA and larger for the AUTO purchase-share outcome. The pooled Heckman shows the same pattern. Wealth lowers fixed entry costs, but larger accounts do not necessarily allocate a larger initial share once they enter.

Risk preference raises first adoption for both products, by 0.7 and 0.8 percentage points, and has no reliable initial-share effect. Female investors show the two-margin pattern on a smaller scale. They are less likely to adopt RA first, but conditional on first adoption they allocate more at entry. The female coefficients are -0.007 in the RA first stage and $+0.026$ in the RA outcome stage, so the entry gap is modest while the initial RA-share difference is about 17 percent of the conditional RA-share mean. Whatever raises the threshold to a woman's first adoption does not carry into a more tentative initial commitment, so a single adoption dummy would record the deterrent and miss the higher starting share. This is a stock-versus-flow sign difference of the same kind we document for fund experience below, and it runs the other way, since the pooled conditional share is lower for women in every panel, at -0.013 to -0.020, while the share at first adoption is higher. Anchoring and extrapolation take

opposite signs between the pooled and first-adoption designs too. We flag these without interpreting them; only the concentration and confirmation patterns are backed by a mechanism test.

Pre-window fund experience reduces first adoption of both products. The coefficient is -0.190 for RA and -0.137 for AUTO, both highly significant. These are the largest first-stage magnitudes in the table, and they are large relative to the first-adoption rates, so the pre-entry experience gradient is the dominant economic pattern in the specification. No other coefficient shows the stock-versus-flow distinction as clearly, because no other reverses the positive pooled participation coefficient of Table 2 so sharply, and that opposition is not a contradiction. The pooled coefficient is positive because experience acts on the stock. Experienced investors who do eventually adopt tend to continue and survive in the RA-active pool, so they accumulate as a disproportionate share of the user-months any monthly snapshot records. The first-adoption coefficient is negative because experience acts on the flow. A larger stock of fund-picking experience gives the investor a better self-directed outside option, so whatever she gives up by ceding fund selection is worth more to her, and she is slower to cross the entry threshold for the first time. The same trait therefore deters the first step and characterizes the surviving pool, which is exactly the lifecycle heterogeneity a single use indicator cannot reveal. Measured selection ability points the same way at entry, lowering first RA adoption by 1.3 percentage points, which is consistent with an outside option that is worth keeping.

Selection on unobservables also changes sign at the entry margin. The first-adoption inverse Mills coefficient is -0.081 for RA and -0.090 for AUTO, opposite to the positive Mills coefficients in the pooled Heckman. Neither is statistically significant, with t-statistics of -0.989 and -0.827, so we read the sign as no more than suggestive. Taken at face value, it points the same way as the rest of the lifecycle evidence. The marginal first adopter, drawn in by something the covariates do not capture, commits a smaller initial share than the average first adopter, which is a cautious-experimenter pattern at entry. The positive pooled coefficient, by contrast, says that among ongoing users the marginal adopter on unobservables delegates more. An investor who starts cautiously and then deepens delegation as comfort accumulates would generate exactly this sign difference between the first-adoption and active-user panels. The insignificance keeps this a reading we offer rather than one we rest on, but it is consistent with the dynamic-delegation interpretation of Section 6.3, not at odds with it.

The first-adoption behavioral signs do more than echo the main Heckman. For the two focal active-opinion biases they remove a reading the pooled estimates cannot. The pooled positive-share coefficient for underdiversification is, on its own, open to a survivorship interpretation, in which concentrated investors who happen to remain in the RA-active pool delegate more and a monthly snapshot simply records survivorship, not awareness. The first-adoption estimates are not open to that reading. Underdiversification lowers first RA adoption and raises the initial RA share, with coefficients of -0.005 and $+0.045$. The positive-share effect, about 30 percent of the conditional RA-share mean, is measured at the investor's first delegation decision, before any post-entry survivorship can operate. The negative-then-positive awareness signature is therefore present at the moment of entry itself, which is the form of the pattern least exposed to composition effects. Confirmation runs the other way at the same moment, raising first RA adoption and lowering the initial RA share, $+0.030$ and -0.040, so the entry effect is about 39 percent of the RA first-adoption rate and the intensity effect about 26 percent of the conditional RA-share mean.

Turnover fits the pattern only partly, and the way it falls short is informative. It lowers first RA adoption, though only at the ten percent level, and has no statistically significant initial RA-share effect, so its entry deterrence replicates while its positive-share counterpart does not appear at first use. This is not an inconsistency with the pooled result. It is a statement about timing. The concentration response is present at the first delegation decision, so whatever produces it is already operating at entry. The turnover pattern is instead a post-entry phenomenon, one that emerges among active users in the pooled Heckman and then continues as the tenure decline in Section 6.1. The two awareness biases therefore correct at different speeds, immediate for underdiversification and gradual for turnover, a distinction the first-adoption design isolates and a single use indicator would hide. Among the remaining biases, attention and herding lower first adoption, consistent with the cognitive-constraint reading, while anchoring lowers both first adoption and initial RA share in this pre-window design, which is one reason we do not force every bias into a single mechanism. The strongest first-entry patterns are concentrated in the same variables that matter in the pooled Heckman, with underdiversification and confirmation producing opposite signs while the attention-related constraints deter entry.

The comparison with first AUTO adoption bears on the same question the ladder design takes up later, namely whether an investor who stays away is avoiding the algorithm or avoiding giving up the choice of which funds to hold. Some predictors move RA and AUTO in the same way at the entry stage, which indicates common product discovery or shared platform channels. Wealth, risk preference, fund experience, anchoring, turnover, attention, underdiversification, confirmation, and herding all have similar first-stage signs for both products. The outcome-stage signs are where the two products part. Turnover lowers AUTO intensity by 4.4 percentage points but has no reliable effect on RA initial intensity. Confirmation raises AUTO first adoption and has a positive but insignificant AUTO-share coefficient, whereas it lowers the initial RA share. Female investors are no less likely to start an automatic plan, though they are less likely to start RA. Execution automation and full delegation are therefore already distinct at the entry point, not merely among active users.

5.4. Competing Accounts of the Baseline Pattern

What the baseline isolates is a puzzle rather than an answer. The investors with the most correctable fund-management mistakes are the least likely to adopt the robo-advisor, yet those who do adopt delegate the largest share of their fund wealth to it. The same trait deters at entry and attracts at intensity, and the pattern is selective: it appears for the biases tied to how a portfolio is run, underdiversification and turnover, in the opposite arrangement for confirmation, and not at all for attention, gambling, and herding, which move both margins together. The selection correction sharpens the puzzle instead of dissolving it, since the positive Mills coefficient says that the investors drawn in by something the covariates miss go on to delegate more once inside.

Two kinds of account can produce that structure. Under the first kind the pattern is not a puzzle at all, because the positive share coefficient measures something other than a decision to commit. Five such readings are available, and each can be settled with evidence already in hand, so we settle them before going further. The first is survivorship. If concentrated investors who happen to remain in the RA-active pool are the ones delegating heavily, the coefficient records who stays rather than who commits. Measured at the investor's first delegation decision, where no post-entry survivorship can

operate, concentration lowers first RA adoption and raises the initial delegated share by 4.5 percentage points, about 30 percent of the conditional mean (Section 5.3). The second is arithmetic. The delegated share divides RA assets by RA plus mutual fund assets, so concentrated investors could show a high share simply by holding a small mutual fund position. Replacing the denominator with total financial assets leaves concentration at $+0.017$; dropping the denominator altogether and using the log level of RA assets leaves it at $+0.204$; and stripping low-risk RA strategies and money market funds from both numerator and denominator leaves it at $+0.052$ (Section 4.5).

The third is product risk. RA might simply be the conservative option, so that investors with an appetite for risk stay out and those without it commit deeply, with concentration standing in for risk appetite. Risk preference is measured and enters the specification directly, and it raises both margins, by 1.9 percentage points on adoption and 1.2 percentage points on the conditional share in the biases panel. That is what a characteristic that does not separate the margins looks like, and it is not what concentration does. The fourth is cost of entry. Adoption might be a wealth problem, with concentration proxying for the resources that clear it. Wealth is also measured, and it behaves as the cost story predicts while concentration does not. In the biases panel, where the behavioral measures enter alongside it, a one-standard-deviation increase in log total financial assets raises adoption by 7.9 percentage points and leaves the conditional share statistically unchanged, at -0.010. The product also carries no minimum, so the entry cost in this setting is discovery rather than money. The fifth is that the behavioral variables proxy for one another, so that the intensive-margin relations run through correlates rather than through the characteristics named. The joint specification of Online Appendix Table A.3, which enters demographics, skills, and all ten behavioral measures together, leaves the underdiversification, turnover, and confirmation sign pairs intact.

The second kind of account treats the pattern as real and offers an economic reason for it. Three do so. One is awareness: the investors who would gain most have not yet seen what delegation would show them, and exposure changes that. A second is authority: investors who stay away are refusing to give up the choice of which funds to hold rather than refusing the machine, so the characteristics that predict wanting an algorithm to act are not the ones that predict letting it choose. A third is dynamic: delegation may be a transitional arrangement that experienced investors unwind. Against all three stands a fourth possibility, that none of them is doing any work and recent returns are the omitted state variable behind the whole pattern. Each of the four implies something the baseline does not report, and Section 6 takes them in turn.

The tests are chosen so that each puts one account under pressure on evidence the baseline cannot supply. An awareness reading implies post-adoption trajectories that decline for the biases delegation makes visible, decline most for the initially worst, and steepen with attention, while the comparison biases stay flat. A generic-automation reading implies that the same characteristics predict AUTO and RA alike, which the within-investor-month ladder can refute. A static-delegation reading implies no tenure gradient in the delegated share. And a performance reading implies that recent returns absorb the characteristic gradients. What these tests do is make the baseline pattern interpretable rather than merely visible.

6. Three Mechanisms

The baseline shows that adoption and delegated share have different determinants in this sample. It does not show why. This section tests three mechanisms, each on evidence the baseline cannot supply, and each answering a different one of the questions the demand for delegated advice raises. Selective awareness answers why entry is slow for the investors who would gain most, and it is tested on post-adoption trajectories rather than on the cross-section that produced the pattern. Decision authority answers which half of the robo-advising bundle investors resist, and therefore what limits commitment, and it is tested on the within-investor-month comparison of the three products. Reclaiming control answers whether delegation lasts, and it is tested by following adopters month by month as their time with the product accumulates. The three are not rival explanations of one fact; they are answers to different questions, which is why each requires its own design. Recent performance is the one alternative that could in principle displace all three at once, and the section ends with it.

6.1. Selective Awareness

Selective awareness is the hypothesis that an investor does not perceive what her own fund-management mistakes cost her until she has a diversified, rule-based portfolio to see them against, and that adopting the robo-advisor is what supplies the comparison. It is an account of why the investors with the most to gain are the slowest to enter, and of what should happen to their mistakes once they do.

The design that tests it tracks the behavioral measures over RA tenure. Monthly bias measures are noisy, so the headline outcomes y_{it} are rolling six-month trailing means of underdiversification, turnover, and the comparison biases. The specifications include investor and month fixed effects with two-way clustered standard errors. The behavioral measures need fund-trading histories to construct, so this design runs on the subsample of investors with constructible biases: 1,321,291 adopter-months in the within-adopter panel and 1,847,939 investor-months once matched never-adopters are added, instead of the full 4.9-million investor-month panel. A matched difference-in-differences version assigns pseudo-adoption dates to never-adopters, so that calendar-time exposure is comparable across treated and control investors. The never-adopter benchmark then enters through an interaction,

$$y_{it} = \rho\,(\text{adopter}_i \cdot \tau_{it}) + \text{controls}_{it} + \mu_i + \lambda_t + \varepsilon_{it}, \tag{12}$$

where τ_{it} is months since first RA use, capped at 24, μ_i and λ_t are investor and month fixed effects, and the adopter indicator and the calendar component common to both groups are absorbed by the fixed effects, so only the interaction is identified. The design is descriptive. What it asks is whether the biases behind the selective-awareness interpretation decline after RA exposure.

The baseline cannot settle that reading. The negative-then-positive sign pair for underdiversification and turnover is what awareness predicts, but a static cross-section cannot show the learning the reading rests on, since it pools investors at every tenure into one snapshot. To see learning we have to watch the mistakes themselves after the investor adopts, which moves the analysis onto event time, months since first RA use, in place of the calendar cross-sections of the baseline. If exposure to a diversified, rule-based portfolio is what makes the cost of these mistakes salient, then the mis-

takes should recede after adoption, and recede most for the investors who were worst at the start and who look at the platform most. The comparison biases, which no mechanism says exposure corrects, should stay flat.

Before reading those trajectories it helps to separate the levels of the argument, because selective awareness operates at two of them at once. At the level of overall demand for delegation, the paper advances three mechanisms: selective awareness of correctable mistakes, decision authority as the substantive content of the delegation step, and reclaiming control over tenure. The behavioral evidence in this subsection sits inside the first. It resolves into three sign signatures, which sit inside the behavioral channel and are not independent paper-level claims. The first is awareness. Underdiversification and turnover are mistakes RA can plausibly discipline, but investors with more severe mistakes may not perceive the cost before adopting. They enter less, and the ones who do enter have crossed the awareness hurdle, which is why they delegate heavily from the first month and why their mistakes then recede with exposure. The awareness signature has a rational benchmark. Investors who optimally learn about the assets they already hold can find concentration optimal rather than mistaken, so its cost need never become salient to them (Van Nieuwerburgh & Veldkamp, 2010), and investors misperceive even their own past returns (Glaser & Weber, 2007).[8] The second is attachment. Confirmation bias can make trial cheap if RA is added as a separate sleeve, but it can reduce commitment because the existing validated positions remain attractive. The third is cognitive constraint. Attention, gambling, and herding biases are not necessarily relaxed by RA, so they can reduce both margins. The differential sign pattern across these three signatures is what distinguishes the behavioral evidence from a uniform account in which all biases respond to RA in the same way.

Table 5 tests the post-adoption prediction directly. It tracks bias trajectories over RA tenure, using rolling six-month trailing means in Panel A and raw monthly bias measures in Panel B. The table splits the biases into the two focal active-opinion measures, underdiversification and turnover, and a set of comparison biases that the framework predicts should not decline with tenure. This split is a prediction about post-adoption trajectories, not about cross-margin signs. Confirmation, for example, is grouped with the comparison set here because the attachment mechanism implies no post-adoption decline, even though confirmation has a genuine and opposite-signed cross-margin signature in the baseline Heckman. The cognitive-constraint biases, attention, gambling, and herding, are reported in Online Appendix Table A.11 and are likewise predicted to be flat. The design tests something stronger than a negative average tenure slope. If RA exposure makes correctable active-opinion mistakes more salient, the decline should be concentrated among high-pre-adoption-bias adopters and amplified by engagement. That is the pattern in the table. In the smoothed panel, underdiversification and turnover decline with RA tenure, at -0.0013 and -0.0029 per month, and the declines are much larger for high-pre-adoption-bias adopters. Over a twelve-month tenure difference, the high-underdiversification interaction implies an additional decline of roughly 0.13 in the concentration measure, while the high-turnover interaction implies an additional decline of about 0.19 in turnover. The active-opinion-taker indicator, defined as high underdiversification or high turnover, also predicts stronger declines in both focal biases. Against pooled standard deviations of 0.318 for concentration

[8] The broader pattern in which the least skilled are least aware of the deficit is documented in psychology (Kruger & Dunning, 1999).

and 0.390 for turnover, these are movements of roughly 0.4 and 0.5 of a standard deviation for the high-bias groups. For turnover this tenure decline also completes the pattern left open at first adoption, where turnover deterred entry but did not yet raise the initial delegated share (Section 5.3). Its correction surfaces among active users and then deepens with tenure, the gradual counterpart to the immediate underdiversification response visible at the first delegation decision.

The matched difference-in-differences rows tell the same story. The Adopter-by-tenure term is negative for underdiversification and turnover, at -0.0008 and -0.0018, and the high-bias and active-opinion-taking interactions remain negative and significant. Engagement is the auxiliary variable that discriminates best among the readings on offer, because the engagement interaction, not the level decline, separates an attention-mediated channel from a mechanical one. A simple negative tenure slope could come from mean reversion or a maturing portfolio, and neither of those predicts a steeper decline for investors who look at the platform more. A salience channel does predict exactly that, since the correction works through the investor seeing the algorithmic portfolio, and seeing requires attention. Login-based attention is a meaningful object in its own right, since portfolio attention is selective, and investors look less after bad news (Sicherman et al., 2016). The Adopter-by-tenure-by-engagement triple interaction is negative and significant for both focal biases in the smoothed panel, at -0.0002 ($t = -3.07$) for underdiversification and -0.0006 ($t = -6.68$) for turnover, and it sharpens in the raw monthly panel to -0.0003 ($t = -4.90$) and -0.0007 ($t = -8.33$).

No comparison bias shows the same engagement pattern, which is the contrast the reading requires, since a general post-adoption drift would dose every bias and does not. Anchoring and extrapolation are indistinguishable from zero on the engagement triple in both panels. Confirmation is the exception within the comparison set, carrying a small positive coefficient of $+0.0003$ in the smoothed panel, so the interpretation is comparative rather than absolute. The focal active-opinion biases display stronger and more consistent post-adoption discipline than the comparison set, and the one comparison bias that moves does so in the wrong direction for a blanket-improvement story. The same holds in Online Appendix Table A.11, where attention, the disposition effect, gambling, and herding all carry positive engagement triples, so the engagement dose is negative for the focal pair and for nothing else.

The comparison biases are not, however, flat in tenure itself, and we do not claim they are. Anchoring rises with tenure, at $+0.0009$ within adopters and $+0.0015$ in the matched design, extrapolation rises in the matched design, and confirmation declines, at -0.0005 smoothed and -0.0022 raw. Their level trajectories therefore do not move in a common direction. That is why the falsification rests on the engagement interaction and not on the tenure slope. A level trend can come from mean reversion, portfolio maturation, or the mechanical effect of moving wealth into a sleeve the investor does not trade, and none of those predicts a steeper decline for the investors who look most. One further caveat belongs here. In the raw monthly panel the simple tenure term and the adopter-by-tenure term for underdiversification lose significance, so the average adopter's decline is visible in the smoothed series and only in subsamples month by month. The heterogeneity and engagement interactions, where the mechanism content sits, survive the removal of the smoothing window in both focal biases. Figure 2 isolates this engagement-dosed falsification test. It plots the negative triple interactions for underdiversification and turnover against the flatter or wrong-signed comparison responses, the visual counterpart of the post-adoption implication of selective awareness.

These trajectories connect the demand side to the effects literature. Prior work shows that robo-advising can improve diversification and reduce mistakes conditional on use (Bianchi & Brière, 2026; Hao et al., 2022; Reher & Sokolinski, 2024; Rossi & Utkus, 2024). Here the same behavioral variables shape voluntary demand and post-adoption behavior. Investors with correctable mistakes are not simply the first to enter; in the baseline Heckman they are less likely to adopt. But conditional on adoption, their subsequent bias trajectories are consistent with learning by doing, or with exposure to rule-based diversified management.

This shifts the corrective-effects literature from an effects question to a selection question. If RA reduces concentration or trading mistakes among users, the relevance for aggregate gains depends on whether the investors with those mistakes enter voluntarily and how much wealth they delegate once inside. The negative adoption coefficients for underdiversification and turnover mark the demand-side bottleneck. The positive conditional-share coefficients and the post-adoption declines locate the margin on which correction appears once the bottleneck is crossed. The paper therefore complements the effects literature by asking who exposes herself to the correction in the first place. This is the resolution of the puzzle the baseline posed. The investors who would gain most avoid RA because they have not yet seen what it would show them, and they delegate hardest, then correct fastest, once they have.

For all this, the taxonomy remains interpretive. The focal-bias declines are consistent with learning after exposure, but they do not prove that investors consciously recognize the cost of concentration or turnover, and mean reversion, concurrent attention shocks, or changing portfolio needs can all contribute to post-adoption movement. What the table establishes is comparative. The declines are stronger for the focal active-opinion biases, stronger among high-bias adopters, and more tightly linked to engagement than the comparison set, a pattern less consistent with a purely mechanical post-adoption trend than with the dynamic implication of selective awareness.

6.2. Decision Authority versus Automation

Decision authority is the right to choose which funds to hold and to place the trades that implement the choice. Robo-advising takes that right; the automatic plan does not, since it automates only the timing of purchases and leaves selling and rebalancing with the investor. The two products therefore differ in which decision the investor gives up, while sharing the platform, the interface, and the execution machinery, which is what makes the comparison in this subsection possible.

The authority ladder separates the two features. It asks whether an investor who stays away from RA is avoiding automation or avoiding handing the choice of funds to the algorithm. The design stacks the DIRECT, AUTO, and RA product rows within each investor-month and encodes the ladder with two cumulative indicators: $c_T = 1$ for AUTO and RA, and $c_S = 1$ for RA only. Investor-month fixed effects absorb wealth, market conditions, a general affinity for automation, and everything else common to the three product rows. Stacking product $k \in \{\text{DIRECT}, \text{AUTO}, \text{RA}\}$ within investor-month it, the estimating equation is

$$\text{use}_{i,k,t} = \alpha_{it} + \gamma_T \left(X_i \cdot c_k^T\right) + \gamma_S \left(X_i \cdot c_k^S\right) + (\text{baseline} \times c_k^T, c_k^S) + \varepsilon_{i,k,t}, \tag{13}$$

with $c_k^T = \mathbf{1}\{k \in \{\text{AUTO}, \text{RA}\}\}$, $c_k^S = \mathbf{1}\{k = \text{RA}\}$, and α_{it} the investor-month fixed effect. The coefficients on $X \times c_T$ then identify the differential associated with automated execution relative to direct investing. The coefficients on $X \times c_S$ identify the further differential associated with RA's full delegation of fund-selection authority. A Wald test of $\gamma_T = \gamma_S$ asks directly whether a characteristic moves the two steps alike. These are within-investor-month product-use differentials, not literal transition hazards.

The account of the intensive margin left open in Section 5.4 is an authority cost, under which investors with more to lose from ceding fund selection, the experienced ones, delegate a smaller share. That account assumes what investors hesitate over is giving up the choice of funds. A simpler alternative is that the investors who use RA simply like financial automation. If that were so, the same characteristics would predict the use of any automated product, and authority would be doing none of the work. The baseline cannot separate these two readings, because it observes RA alone, and RA bundles automation together with delegation.

Table 4 separates the two readings on 672,228 stacked observations. Because the comparison is made inside the investor-month, it has no time dimension: it asks how the same person, in the same month, ranks the three products, not how anyone changes over time. The language of steps is an authority-ladder reading of within-investor-month product differentials and should not be mistaken for a sequential transition model. That the two products are not interchangeable is what makes the ladder necessary rather than decorative. Estimated on its own, AUTO has different determinants from RA: underdiversification raises AUTO intensity by 8.8 percentage points against 2.0 for RA, and turnover is negative on both AUTO margins where it is positive on the RA share (Online Appendix Table A.6).

The central result is that the covariates do not move the two steps in parallel, and for several of them the two steps carry opposite signs. Platform engagement is the sharpest case. Pre-use login frequency lowers the automation step by 4.2 percentage points and raises the full-delegation step by 1.6, and the equality of the two is rejected decisively ($\chi^2 = 29.38$, $p < 0.001$). Total turnover does the same with the signs exchanged, raising the automation step by 3.5 percentage points and lowering the full-delegation step by 3.2 ($\chi^2 = 31.80$, $p < 0.001$). Underdiversification is the largest single split, at $+5.8$ and -3.3 percentage points ($\chi^2 = 69.60$). Anchoring runs -3.8 and $+4.8$ ($\chi^2 = 57.79$), and attention bias -4.2 and $+2.6$ ($\chi^2 = 29.63$). Figure 1 plots the composite coefficients with their confidence intervals and the equality tests, and what the figure shows is that the covariates do not shift the DIRECT-to-AUTO and AUTO-to-RA steps together. These are percentage-point differentials in product use within the investor-month, and they are large against the rates at which the two products are held: an automatic plan appears in 32.9 percent of investor-months and RA in 32.1 percent (Section 4.1). A split of $+5.8$ against -3.3 therefore moves each step by roughly a tenth of its own base. A generic-automation account cannot produce any of these opposite-sign splits. Under such an account the second step would carry no differential at all, since an investor drawn to automation as such has no reason to treat the more automated product differently once the first step is taken. What we find instead is a second step that is significantly non-zero and opposite in sign to the first. The two product differentials are therefore economically distinct choices, not two labels on a single automation preference.

Where the two steps take opposite signs, the direction of the split is clear, and it is what an authority

cost predicts. Among the variables whose steps differ, those that push toward automation and away from delegation are the ones measuring how much active opinion the investor holds about which funds to own: total turnover and portfolio concentration. The composite authority-bias index points the same way at the delegation step, at -0.023 ($t = -3.90$) against an automation coefficient of essentially zero, with equality rejected at the five percent level ($\chi^2 = 4.10$). Scheduling automates work these investors are content to keep; full delegation would hand off the allocation decision they have views about, so they tilt toward the first relative to the second. The variables that take opposite signs the other way measure a binding constraint on attending to that decision: engagement, attention bias, and anchoring. For those investors the automatic plan solves nothing, because their problem is not the mechanics of placing trades, while full delegation solves the allocation problem outright.

This reading covers the rows with opposite signs, and five of the twelve disaggregated rows sit outside it. The sharpest of those is speculative turnover, which is negative on both steps, at -0.031 and -0.013, and keeps both signs at every horizon the appendix reports, so these investors decline to hand the decision to either product. The others, confirmation, gambling, herding, and the disposition effect, load on one step or on neither, and we state the split for what it is, the pattern among the opposite-sign rows rather than a law governing every covariate. The falsification does not depend on it. Five rows carry opposite signs with both step coefficients individually significant, namely total turnover, portfolio concentration, anchoring, attention, and engagement, and equality is rejected at $p < 0.001$ in every one. Whatever generates that, it is not a single taste for automation.

Measured skill is the row that departs most cleanly from what an outside-option reading would predict. The skill composite is unrelated to the automation step ($\gamma_T = 0.004$, $t = 0.75$) and positive on the full-delegation step ($\gamma_S = 0.012$, $t = 1.79$), though the Wald test does not reject equality ($\chi^2 = 0.54$). Selection ability is sharper, at -0.006 on automation and $+0.014$ on delegation ($\chi^2 = 3.65$, $p = 0.056$), and it strengthens at longer pre-use windows, where the skill composite reaches $+0.016$ on the delegation step at the five percent level and turns negative on the automation step (Online Appendix Table A.9). Investors who have been picking better funds are not the ones buying mere scheduling; if anything they are the ones willing to hand the selection over. Two qualifications keep this modest. The Wald test never rejects equality for the skill composite at five percent at any horizon, so the two step coefficients are barely distinguishable. And the other skill measure, timing ability, runs the other way, loading weakly on automation ($+0.009$, $t = 1.83$) and not at all on delegation. What we take from the row is therefore narrow. Measured ability does not predict a taste for mere scheduling, so the authority cost in this setting attaches to opinion intensity, not to measured ability.

The composite indices carry labels that predate the estimates. The skill index averages selection and timing ability; the authority-bias index averages confirmation, both turnover measures, herding, and the disposition effect; the placebo index averages underdiversification, gambling, anchoring, attention, and extrapolation. The placebo label records what the authority mechanism was expected to load on, not a claim that its components are inert, and two of them plainly are not. So constructed, the placebo index is indistinguishable from zero at both steps (-0.007 and $+0.003$) with no cross-step difference ($\chi^2 = 0.74$), even though underdiversification and attention inside it carry large opposite-signed step coefficients. The index nets them out, which demonstrates the paper's general point at the level of a single variable: aggregating behavioral measures averages opposite-signed demand signa-

tures toward zero, and only the disaggregated estimates recover them.

One apparent tension deserves stating plainly, because it is a feature of the design rather than a contradiction. Underdiversification raises the conditional RA share in the pooled Heckman by 2.0 percentage points, yet tilts investors toward AUTO and away from RA in the ladder. The two designs measure different objects. The pooled Heckman asks how much fund wealth an adopter places under RA; the ladder asks how an investor's use of the three products differs within a month. The AUTO Heckman reconciles them, since underdiversification raises AUTO intensity by 8.8 percentage points against 2.0 for RA, so concentrated investors delegate meaningfully to RA when they adopt while using the fund-choice-preserving product relatively more. Both statements can hold, and they do.

Behind these coefficients sit two recognizable investors. One has strong views about which funds to own and can value automated execution because it reduces the effort of placing trades while preserving her right to choose, yet need not want RA to take the choice of funds as well. The other is constrained by attention rather than conviction, and she runs the opposite way, caring little for scheduled execution but valuing a product that solves the allocation problem more completely. The table thus separates two margins that discussions of fintech adoption usually bundle, automation of the task and delegation of the decision, and the sign differences are evidence about the latter, not about digital product use in general.

What the ladder establishes is that RA demand is not a generic preference for digital automation. What separates the two products is that RA takes the choice of funds out of the investor's hands, and Section 2 places that finding against the experimental work on decision rights and the evidence on algorithm aversion. We do not claim that fund-selection authority is the only difference between the two products. AUTO and robo-advising are matched on platform, account interface, liquidity, and execution infrastructure, and the within-investor-month design absorbs everything common to the two product rows. The products may still differ in branding, perceived sophistication, default framing, or formal advisory status. Authority is the most economically interpretable of these differences and the one our covariates are designed to track. We therefore read the within-investor-month tilt as evidence about authority, while acknowledging that these other product attributes are bundled with it rather than fully stripped away.

6.3. Reclaiming Control Over Tenure

Reclaiming control is the withdrawal of authority already handed over: an adopter who lowers her delegated share as tenure accumulates takes back part of the fund-selection decision she transferred at entry. It is the dynamic counterpart of the authority cost, and it treats delegation as a state an investor can leave by degrees rather than an absorbing one. The authority result just reported is static, describing how an investor ranks three products in a single month; this subsection adds the time dimension the ladder cannot supply.

If experience raises the value of keeping fund-selection authority, then adoption should not be the end of the matter. An experienced investor who delegates should want that authority back as she relearns that fund selection is something she can do herself, so her delegated share should fall the longer she has held RA. A single cross-section cannot see this, because it pools investors at every tenure into one snapshot. The relevant clock is therefore no longer calendar time but time since adoption, and

the test has to follow the delegated share over RA tenure and ask whether the fall is steeper for the investors the authority-cost reading says have the most to reclaim.

The test estimates the RA share among RA-active investor-months as a function of pre-RA fund experience and the interaction of that experience with RA tenure:

$$\text{rsh}_{it} = \beta \, \text{exp}_i + \delta_1 \left(\text{exp}_i \cdot \widetilde{\tau}_{it}\right) + \delta_2 \left(\text{exp}_i \cdot \widetilde{\tau}_{it}^2\right) + \phi_1 \widetilde{\tau}_{it} + \phi_2 \widetilde{\tau}_{it}^2 + \text{controls}_{it} + \varepsilon_{it}. \tag{14}$$

Pre-RA experience exp_i is measured over the six months before first RA use and standardized across eligible adopters. Six months is long enough to measure the investor's habits stably and short enough to stay clear of behavior that anticipates adoption; Online Appendix Table A.9 varies the analogous window in the authority design. Tenure $\widetilde{\tau}_{it}$ is the cumulative number of RA-active months, demeaned within calendar month. The sign of δ_1 answers the design's question, namely whether investors who had more fund-market experience before adopting cut their delegated share faster as tenure accumulates. Equation (14) is estimated by Fama–MacBeth across monthly cross-sections with Newey–West standard errors, and re-estimated with a fractional-logit link as a functional-form check for the bounded share. Because each adopter needs six months of pre-RA history, investors already using RA at the start of the window cannot enter, and the design runs on 1,097,096 of the 1,571,475 RA-active investor-months across 44 monthly cross-sections (Table 8).

The evidence is consistent with this reclaiming-control reading, with one important precision. The Fama–MacBeth estimates describe within-month cross-sectional tenure gradients among active RA adopters, not a pure within-investor time derivative. Longer-tenure active adopter-months have lower RA shares on average, in both the OLS and the fractional-logit specifications. The linear OLS tenure coefficient is -0.028, the quadratic OLS specification gives -0.048, and the fractional-logit specification gives -0.330 on the logit scale. Pre-RA experience is negative in levels, about -0.196 in the OLS specifications. The focal interaction between pre-RA experience and tenure is also negative: -0.023 in the linear OLS specification, -0.043 in the quadratic OLS specification, and -0.304 in the fractional-logit specification. The quadratic interaction terms are negative but insignificant throughout, so the decline is adequately described as linear in tenure. Online Appendix Table A.10 preserves the sign and size under alternative adopter-sample restrictions.

Figure 3 translates the linear interaction into the tenure slope of RA share at different levels of pre-RA fund experience. The slope is close to zero for the least-experienced adopters, at -0.005 one standard deviation below the mean, and grows more negative as pre-adoption experience rises, reaching -0.051 one standard deviation above it. The slope crosses zero at about 1.2 standard deviations below mean experience, so for all but the least experienced adopters longer tenure goes with a lower delegated share. That gradient gives the interaction its economic content.

Economic size bears on whether RA is a durable delegation product. The dependent variable is a share bounded between zero and one, and the conditional RA-share mean in the main Heckman is 0.151. Even allowing for tenure being demeaned within calendar month, the OLS magnitudes imply that relative tenure differences and pre-adoption experience differences move a large fraction of the average delegated share. The negative interaction is therefore not merely a sign test. Within a month, the tenure-share gradient is more negative for adopters with more pre-RA fund experience. We do not overstate the precision. The focal interaction is significant at the ten percent level in the headline linear

specification, at five percent in the quadratic and fractional-logit columns, and at five or ten percent across the three appendix variants, so the steepening with experience is well signed but not sharply estimated, while the tenure decline at and above mean experience is significant at conventional levels.

This result complements the authority ladder. The ladder speaks to authority transfer at a point in time; the tenure test suggests that investors with stronger pre-adoption self-directed experience hold lower delegated shares at longer RA tenures. This dynamic pattern also disciplines how we read RA use. Adoption is not an absorbing state, and the evidence does not support a simple monotone deepening story. RA can serve as a temporary or partial delegation technology, particularly for investors who retain enough fund-market experience to value control over the self-directed sleeve. The same pattern appears from the self-directed side in the event-study evidence of Section 7.1, where portfolio concentration falls sharply at adoption and then recovers toward its pre-adoption level over the following months, and in Online Appendix Figure A.2, where the within-adopter concentration gap moves from -0.027 at month five to $+0.023$ by month thirty-six while the RA share returns to its adoption-month level. The delegated share declining over tenure and the self-directed sleeve rebuilding over the same horizon are two views of one reallocation back toward self-direction.

Tenure also clarifies the role of experience in the main Heckman. Experience is more than a proxy for sophistication or wealth. In pooled participation it can lower discovery and evaluation frictions. At first adoption and on the conditional-share margin it can also raise the value of retained authority. That pattern fits learning by trading: the longer an investor has been picking funds herself, the better she gets at it, and the more she gives up by handing the job over (Arrow, 1962; Ben-Porath, 1967). The investor can delegate, but she continues to value the authority to manage the self-directed sleeve. Read together with the bias trajectories of Section 6.1, the persistence evidence adds a further implication for what success means for corrective advice. The mistakes the algorithm visibly corrects fade with exposure, fastest among the initially most biased and most engaged adopters, while the adopters with the strongest self-directed capability reduce their delegated share as tenure accumulates and rebuild the self-directed sleeve. Help that visibly corrects a mistake also shows the investor what the correction looks like, and the investors best equipped to act on that lesson take the decision back over time, in the lineage of learning by trading (Linnainmaa, 2011; Nicolosi et al., 2009; Seru et al., 2010). The pattern also qualifies two adjacent accounts of delegation. In rational lifecycle models, delegation absorbs the rising time cost of self-management and should deepen with tenure (Kim et al., 2016), and retirement-plan participants rarely adjust their allocations themselves (Agnew et al., 2003; Sialm et al., 2015); here, voluntary delegation on an open platform unwinds precisely for the investors whose self-directed capability is highest. The direction also parallels a human-factors concern, in which reliance on automation erodes the operator's skill (Parasuraman & Riley, 1997); here exposure builds skill, and disuse follows competence. Part of the value of corrective delegation may therefore be realized precisely when users graduate from it, a possibility that cautions against measuring the success of such products by assets retained alone.

6.4. Alternative Explanations and Concerns

Performance is the leading alternative to a characteristic-based interpretation. If investors simply chase recent returns or abandon underperforming sleeves, then the demographic, skill, and bias co-

efficients might be proxies for performance feedback. This subsection is therefore an alternative-explanation check, not a fourth mechanism alongside the three above. The question is narrow, namely whether recent returns can explain away the two-margin characteristic patterns. If they could, the paper would be less about awareness, authority, and the dynamics of control, and more about return feedback. The answer below is no. Recent returns do move demand, but they leave the characteristic patterns in place, so they are part of the story rather than the whole of it.

The designs that test it put return measures into the RA-share equation and into the first-adoption and lifecycle specifications. The appendix reports single level-return controls for fund and RA performance. The main return table decomposes fund and RA returns into positive and negative components, $\beta^{+}r_{it}^{+} + \beta^{-}r_{it}^{-}$, and examines above-median return indicators. The lifecycle-return table studies first adoption, exit, and continuation separately. Entered as single level measures, fund and RA performance are economically small in the RA-share equation, with coefficients no larger than 0.003 across nine specifications (Online Appendix Tables A.7 and A.8). That is not evidence that returns are irrelevant, but evidence that a level control is the wrong instrument for finding them.

Table 7 recovers what a single level control misses. The small fund coefficients in the single-measure specifications reflect a signed return that blends a gain effect and a loss effect, and once those are separated they sit on the same side and net toward zero. Entering each return as its positive and negative parts, with the fund return in both stages, makes that structure visible. The first thing it shows is that own fund performance pushes adoption and conditional delegation in opposite directions. In the raw decomposition both halves of the lagged fund return lower the probability of RA adoption, by about 2.6 to 3.1 percentage points for the positive part and 2.4 to 2.7 for the negative part per within-month standard deviation. Each is significant at the one percent level, and each is roughly 7 to 10 percent of the participation mean. Larger recent own-portfolio outcomes, whether gains or losses, make an investor less likely to enter. It is therefore the size of recent performance, not its direction, that deters entry. The above-median indicator is flat in the selection equation at every horizon, the same conclusion reached from the other side, since crossing a good-versus-bad threshold does nothing once the magnitude of the move is set aside.

Among adopters the sign goes the other way. Fund return^{+} raises the conditional RA share at all three horizons in the raw decomposition, with coefficients of 0.006 to 0.009, and fund return^{-} is also positive, 0.006 to 0.008. These are roughly 4 to 6 percent of the 0.151 conditional mean per within-month standard deviation, and close to one another in size. That near-symmetry is what discriminates among the rival readings, since each predicts something different about it. A complacency channel, in which good own returns breed comfort with delegating, would load only on fund return^{+}. A disappointment channel, in which poor own returns chase investors into the algorithm, would load only on fund return^{-}. A dispersion channel, in which the magnitude of recent own outcomes moves delegation regardless of sign, is the only one that places both components on the same side. That is what the data show.

The RA return, observed only in the outcome stage, behaves differently across its two halves. RA return^{+} is positive at all three horizons, at the ten percent level after one month and at one percent after three and six, with coefficients of 0.006 to 0.009 comparable to the fund components, so in the gain region own and delegated recent performance raise delegation by similar amounts. RA return^{-}

is positive throughout but reaches significance only at the six-month horizon, so in the loss region it is mainly own performance that moves the share. The excess-return decomposition sharpens the contrast. The fund components stay positive and significant, at 0.008 to 0.012, while the robo excess components are small or negative and the above-median indicators are weak. This points away from a relative-performance race between the two sleeves and toward the salience of the investor's own-side outcomes.

Return feedback therefore operates through magnitude, not sign, on both margins, which places performance in the paper as an alternative channel and a boundary condition rather than as a fourth mechanism. It shifts adoption through the magnitude of recent own returns, and it explains variation in allocation inside active RA states once returns are split into positive and negative components. In both places, though, the effect is modest. The level effects, a few percentage points on adoption and at most about one percentage point on the conditional share, are small next to the wealth, experience, and bias coefficients, and they leave the larger opposite-sign patterns on those characteristics intact. The inverse Mills coefficient in these specifications, between 0.091 and 0.144 and significant throughout, is also close to the main-table value, so adding returns does not absorb the residual selection channel either. Recent performance does not overturn the interpretation that stable investor characteristics and behavioral states separate adoption from conditional delegation.

None of this contradicts the experimental algorithm-aversion evidence. The experiments in Dietvorst et al. (2015) and Dietvorst et al. (2018) isolate reactions to algorithmic errors in a setting where the algorithm's performance signal is sharp and the decision maker can attribute the error to the algorithm. Our return variables are monthly sleeve-level portfolio signals from a field setting, and they combine market exposure, fund composition, rebalancing, investor cash flows, and product mix. A weak or sign-symmetric response to these signals should therefore not be read as rejecting algorithm aversion. It says that, in this panel, monthly realized returns are too noisy and too bundled to be the primary object organizing RA demand, while authority and behavioral state explain more of the stable cross-margin structure.

7. Consequences for the Self-Directed Portfolio

The mechanism tests established what separates the two margins in demand. This section asks what the two margins leave behind in the portfolio the investor still manages herself, and how far the estimates survive alternative samples, estimators, and outcome definitions. Two designs do the work, and both separate extensive- and intensive-margin associations. The event-time figures describe changes around adoption, tracing coefficients π_j on relative-time indicators $\mathbf{1}\{\tau_{it} = j\}$ relative to the month before adoption. The dose-response table relates lagged RA-share bins to self-directed-sleeve outcomes. What they show is whether entry and commitment behave as empirically different margins in the self-directed portfolio, as well as in the demand equations.

One caveat governs everything that follows. The estimates are demand estimates, not causal treatment effects, since adoption is chosen. Adoption is also staggered across investors, and event-study coefficients estimated on staggered adoption need not recover a common treatment path (Roth et al., 2023). We therefore read the event-time and dose-response evidence as descriptions of the adopter

panel, and the test they jointly apply is comparative, whether entry and delegated share leave different marks on the self-directed portfolio, as three distinct mechanisms would imply and a single participation decision would not.

7.1. Event-Time Evidence Around Adoption

The event-study figures speak to the extensive margin, that is, to what happens to the self-directed portfolio around the act of adopting, regardless of how much is delegated. Around adoption spells, self-directed portfolio concentration falls sharply. The spell-level event study in Online Appendix Figure A.1 shows the Herfindahl index falling by about 0.038 at spell start, with a trough 0.057 below the reference four months after adoption. The delegated share jumps at adoption, by 17.6 percentage points, and then declines over the first year, to 8.9 percentage points above the reference by month twelve. Turnover rises in the adoption month, by 0.024 in units of the monthly turnover rate, consistent with mechanical portfolio reallocation at entry, and then falls well below its pre-adoption reference in the following months, to roughly -0.12 by month twelve. The tenure event study shows that some of the initial diversification effect unwinds over time, with concentration recovering gradually over the post-adoption horizon.

Figure 4 makes the timing transparent. In the first-time adoption series, the reference HHI in the month before adoption is 0.403. HHI falls to 0.365 in the adoption month and reaches about 0.342 four to five months after adoption, roughly a 15 percent decline in concentration relative to the pre-adoption reference. A movement of that size is economically meaningful against the household benchmarks, since concentrated retail fund portfolios carry diversification losses of the order documented by Goetzmann and Kumar (2008) and Calvet et al. (2007). The effect then attenuates, reaching 0.358 by month 12, 0.381 by month 18, and 0.403 by month 24, indistinguishable from the reference. The any-adoption series, which includes re-entries, follows the same pattern over its shorter window. The event-study evidence therefore does not show a permanent one-time transformation of the self-directed sleeve. It shows a large entry-associated diversification movement that fades over the following two years. Slow reversion of this kind is what the household rebalancing literature would lead one to expect, since individual investors adjust portfolio weights only partially and with long lags (Calvet et al., 2009).

Adopters who also use the platform's other advisory products hold broader and less heavily traded self-directed sleeves in those months, with 8.93 more funds and concentration lower by 0.065, so delegation does not mechanically displace self-direction (Online Appendix Table A.18). The event-time pattern is therefore consistent with a discipline-associated shift in the self-directed sleeve around RA adoption. The extensive margin of entry is associated with economically meaningful changes in the self-directed portfolio, even before the size of the delegated share is taken into account.

7.2. Dose Response in the Delegated Share

The event study held the size of delegation aside; the dose-response design brings it back. Table 9 relates lagged RA-share bins to self-directed-sleeve outcomes in a two-way fixed-effects panel, with no lagged RA share as the omitted category. If a larger delegated share were simply more of the same entry effect, the sleeve outcomes would deepen monotonically across the bins. They do not. Low and

moderate RA-share bins are associated with lower self-directed concentration, more funds, and lower turnover. Bin 1, with lagged RA share in $(0, 0.20]$ and covering 26.8 percent of panel rows, reduces HHI by 0.216 standard deviations, increases the number of funds by 0.196 standard deviations, and lowers turnover by 0.105 standard deviations. Bin 2 has the same signs, though smaller. The middle bin is transitional. HHI remains slightly lower, but the number of funds falls and turnover rises.

Figure 5 plots the binned coefficients, and the non-monotonicity is the substantive point. The low-dose bins and the near-total-delegation bin describe different residual self-directed sleeves. The very high RA-share bins look different again. Bin 4 raises HHI by 0.117 standard deviations and turnover by 0.299 standard deviations, while reducing the number of funds. Bin 5, where lagged RA share exceeds 0.80, reduces the probability of holding any mutual fund by 33.2 percentage points, raises concentration by 0.396 standard deviations, lowers the number of funds by 0.122 standard deviations, and raises turnover by 0.805 standard deviations among mutual-fund-active observations. The residual self-directed sleeve at very high RA shares is therefore smaller, more concentrated, and more actively traded.

A higher RA share does not simply scale up the extensive-margin adoption association. Moderate delegation is associated with broader and less heavily traded self-directed portfolios, consistent with the robo sleeve absorbing the marginal, lower-conviction positions. Near-complete delegation leaves a smaller, more concentrated, more actively traded residual. That residual is consistent not with random shrinkage but with the investor deliberately keeping outside RA the handful of conviction positions she most wants to manage herself. The intensive margin thus changes the composition of what remains outside RA, which is what an authority-cost framework would lead us to expect, since the cost of ceding control is highest for the positions she has the strongest views about. The alternative reading, that the residual is simply what naive allocation leaves behind (Benartzi & Thaler, 2001), does not fit the dose pattern, because a naive residual would not become more concentrated and more actively traded as the delegated share rises. Online Appendix Table A.15 supports that reading from the placebo side, since the same bins leave gambling bias entirely unmoved and move the disposition effect in only one of the five, so what changes across the dose is largely the composition of the residual portfolio rather than the investor's underlying behavioral type.

7.3. Performance Across the Lifecycle Margins

If recent returns were the state variable driving RA use, they should show up most clearly at the moment of first entry, when the investor is deciding whether to adopt at all. Online Appendix Tables A.12 and A.13 add prior self-directed fund returns to the first-adoption specifications to check this. The coefficients are generally small or inconsistent across horizons and return measures, and they take opposite signs between the pre-window and lag-1 hazard constructions. Before first adoption the investor has no individual RA return history, so the only available return signal is her self-directed fund portfolio, which makes the weak result informative rather than merely null. Investors do not appear to enter RA simply because their recent fund returns are low, or because they are mechanically chasing a recent fund-return state.

Table 10 separates the return story across the adoption lifecycle. Because the fund return enters the selection stage of each design, the table shows how recent own performance moves the probability of

using the robo-advisor at each margin. Section 6.4 established that the return effects are economically small. This table adds that they are also margin-specific. The magnitude of recent returns deters first entry, while the sign of recent returns drives exit and continuation, so performance enters the lifecycle in a structured way instead of as a single return-chasing force. Panel A studies first adoption. In the selection stage, both halves of the lagged fund return lower the probability of a first adoption, by about 0.2 to 0.3 percentage points for the positive part and about 0.6 for the negative part per within-month standard deviation, each significant at the one percent level. Against a first-adoption rate near 3.2 percent these are economically large, with the loss component close to a fifth of the base rate. That 3.2 percent is a third base rate, distinct from both the 1.43 percent panel rate and the 7.7 percent pre-window rate, because this design conditions on the at-risk months of not-yet-adopters and drops post-adoption months; each margin's coefficients should be read against its own denominator. This is why the general performance controls in the appendix look weak at first adoption. A signed return averages a gain effect and a loss effect that point the same way, so it is the magnitude of recent own performance, not a low return on its own, that holds investors back from the first step. The reading is not return-chasing, since larger gains deter entry just as surely as larger losses. The initial-share response in the outcome stage is muted at short horizons and turns positive and significant only at six months, where fund return^{+} and fund return^{-} reach 0.024 and 0.042. Panel B treats exit as a probit-only design and recovers the opposite pattern. Good recent fund performance raises the probability of leaving the robo-advisor, by about 0.1 to 0.3 percentage points across horizons, while fund losses lower it by 0.1 to 0.2. The robo-advisor return matters mainly through a positive lag-one coefficient on RA return^{-}, so a recent month of poor robo performance nudges investors out.

Panel C, continuation, completes the same shape. Good fund returns reduce continuation in the selection stage and fund losses raise it, the same direction as exit. The continuation rate sits near one, so these marginal effects describe the marginal would-be exiter and are best read as signs, not levels. The direction of the exit response is the one reinforcement learning predicts (Choi et al., 2009), where investors reinforce the activity that has just paid off, and the activity a good self-directed return rewards is managing the funds oneself, so success on the outside option pulls the investor out of delegation. That performance moves exit at all distinguishes this setting from the retirement-plan evidence, where accounts are held through default and inertia and respond weakly to returns (Agnew et al., 2003; Samuelson & Zeckhauser, 1988). Inside active allocation the intensive-margin return effect is the strongest of the three margins. Both fund-return components raise the conditional robo-advisor share in Panel C, with coefficients around 0.012 to 0.014, roughly 8 to 9 percent of the conditional RA-share mean. The robo-advisor return contributes through its positive part at every horizon and through its negative part at six months. Return feedback is therefore most visible once investors are active, faint and slow at first entry, and directional at exit and continuation. The lifecycle pattern reinforces the central distinction of the paper. Recent performance bears on whether an investor stays or leaves through its sign, but on how much she delegates through its magnitude, and neither reading displaces the characteristic-based separation of adoption from conditional delegation.

The lifecycle evidence sharpens the interpretation of who uses RA. Initial entry is shaped by pre-adoption wealth, experience, risk preference, gender, and behavioral state. Active RA allocation is shaped by the same characteristics, but also by return salience and accumulated product experience.

Exit and continuation respond to performance in ways a one-sided return-chasing model does not describe well. A paper that treats RA use as a single binary outcome would pool these states and obscure the economic content of each margin.

7.4. Interpretation and Limits

The consequence evidence supports a narrow, margin-specific claim. The event-time figures show that self-directed-sleeve outcomes change around RA entry. The dose-response table shows that those outcomes vary non-monotonically with lagged RA-share bins. The adopter-panel sleeve regressions show that active use of the platform's other advisory products can coexist with substantial self-directed investing without mechanically replacing it. Together these results show that entry and delegated share are empirically different margins, and that the self-directed sleeve does not respond to them in a single monotone way. They are not, however, a basis for quantifying a causal total effect of RA, or for decomposing such an effect into entry and commitment components. Recovering the welfare consequences of delegation would require both a causal design and a mapping from portfolio changes to utility, of the kind the diversification-welfare literature supplies (Calvet et al., 2007; Rossi & Utkus, 2024); we do not attempt it.

7.5. Robustness

The appendix collects the robustness checks, organized around the claims they defend. With one exception, discussed last, none of them changes a sign, a magnitude worth re-reporting, or the reading of any result, and the tables are in the Online Appendix. The cross-margin signs are not an artifact of specification or estimator. The joint Heckman of Online Appendix Table A.3, which enters demographics, skills, and biases together, leaves the underdiversification, turnover, and confirmation sign pairs intact, and the panel control-function design of Online Appendix Table A.5, which replaces Fama–MacBeth with investor fixed effects and two-way clustering, recovers the same intensive-margin signs on both the main and the risky-share outcomes. The one attenuation, fund experience on the conditional share in the joint specification, is already how the main text reads that relation.

Nor are the signs an artifact of the outcome definition, the products' distinctness, the sample, or the windows. The alternative intensive outcomes of Online Appendix Table A.4 preserve the behavioral slopes, and the selection term is positive on the share outcomes and negative on the level outcome, which fits the reading of the unobservable as a willingness to cede control rather than a wealth proxy. The stand-alone AUTO Heckman of Online Appendix Table A.6 confirms that AUTO and RA have distinct determinants, as the authority ladder requires. Online Appendix Table A.9 varies the pre-use horizon and relaxes the sample restriction, and the five equality rejections the main text leans on hold at every reported horizon; the checks we do not lean on, extrapolation and the composite indices, are the ones that vary. Online Appendix Table A.10 re-estimates the reclaiming-control interaction under three adopter-sample restrictions and it stays negative and of similar size throughout. The calibrated-weight re-estimation of Online Appendix Table A.1 preserves the cross-margin sign pairs under population weighting, with underdiversification and turnover at $+0.055$ and $+0.062$ on the conditional share against confirmation's -0.039, so the choice to report unweighted estimates is one of precision, not of behavioral findings. The selection term is the exception: the inverse Mills coeffi-

cient turns negative and significant under weighting, at -0.098 to -0.197 against $+0.158$ to $+0.212$ unweighted. Reweighting changes the composition of the estimation sample toward the population of non-adopters, and the selection term is the coefficient most exposed to that change, since it is about who is drawn in rather than what they do once inside. We therefore read the trust interpretation of Section 5.2 as conditional on the unweighted estimand, and none of the paper's mechanism evidence rests on it. The performance appendices, Online Appendix Tables A.7, A.8, and A.14, confirm that return controls leave the characteristic gradients where the main text reports them. Sample splits and cohorts in Online Appendix Table A.16 and Figure A.3 check the signs across windows, and Online Appendix Tables A.17 and A.18 document the age-bias correlations and the self-directed sleeve used in the interpretation sections.

The exception is the exclusion restriction, and here we report a sensitivity that does move. Online Appendix Table A.2 re-estimates the model with platform tenure in both equations (Arm B) and in neither (Arm C), against the benchmark with tenure in the selection equation only (Arm E). The three arms share one bias block of nine measures, which omits speculative turnover, so the benchmark arm's outcome coefficients differ slightly from Table 2, at $+0.018$ against $+0.020$ for underdiversification; the comparison that matters is across arms within this table. The outcome-stage slopes are not stable across the arms. In the biases specification, the underdiversification coefficient moves from $+0.018$ in the benchmark to -0.042 in Arm B, confirmation from -0.023 to an insignificant $+0.006$, and turnover keeps its sign but loses significance, at $+0.009$ against $+0.039$.

These are not separate movements. Across all fifteen outcome coefficients the table reports, the change from the benchmark is 0.30 times that variable's own selection coefficient, with a correlation of 0.997 and an R^2 of 0.99 through the origin. One number therefore reproduces the entire Arm B coefficient vector from the benchmark vector and the selection equation, which is the algebraic signature of a correction that is not separately identified rather than a set of independent robustness failures. Without a variable excluded from the share equation, the correction term depends only on the probit index. That index is itself a combination of the same regressors, and over the relevant range the term is nearly a straight line in it. Adding the term is therefore close to adding the index again, which shifts every coefficient in the same direction at once. Which coefficients visibly change sign is then determined by whether the benchmark coefficient is smaller in absolute value than 0.30 times that variable's own selection coefficient. For underdiversification that threshold is $0.30 \times 0.172 = 0.052$, above the benchmark coefficient of 0.018, so the sign changes; for turnover it is $0.30 \times 0.084 = 0.025$, below the benchmark coefficient of 0.039, so it does not.

The table's own diagnostics say the same thing in other units. The R^2 from regressing the inverse Mills ratio on the outcome regressors is 0.995 in Arms B and C against 0.988 in the benchmark, the condition number of the outcome design matrix rises from 96 to 138, and the implied $\hat{\rho}$ reaches 3.5 to 4.2, far outside the admissible correlation range. The benchmark is not itself immune. Its implied $\hat{\rho}$ in the biases specification is 1.381, and the joint specification of Online Appendix Table A.3 also exceeds one, which is why we report $\hat{\lambda}$ rather than $\hat{\rho}$ as the selection statistic throughout. The difference between 1.4 and 4.2 is nevertheless the difference between a correction that is doing too much work and one that is not separately identified. The selection term in those arms absorbs variation it cannot legitimately claim, and the behavioral slopes move because of it. What the exercise establishes is

that the exclusion restriction is doing real work in the second stage, and that the selection-corrected outcome slopes should be read as conditional on it. It does not establish that those slopes are robust to abandoning it, and we do not claim that they are. Section 4.5 sets out what does not depend on the exclusion. The adoption results do not: the selection equation is the same whether tenure is included or dropped. Nor does the authority ladder, which has no first stage and so no excluded variable. First adoption and the panel control function keep the exclusion but change the sample and the variation used, which makes them weaker checks than the ladder.

8. Conclusion

Robo-advising delegates an investment decision rather than merely digitizing access to markets, and that distinction changes how its demand should be measured. We use Alipay panel data from September 2020 through September 2024, in a setting where the algorithm executes what it recommends. A client who refuses the product and a client who fails to follow through therefore cannot be confused. That lets us separate three decisions a household makes about delegated advice: whether to adopt, how much fund wealth to place under the algorithm, and whether to keep it there, the last two being how far the household commits. In this setting they respond to different characteristics, and several characteristics take opposite signs across them, so a binary use regression pools objects a delegated product keeps distinct.

The argument runs on a small number of headline estimates. A one-standard-deviation increase in portfolio concentration lowers adoption by 6.5 percentage points and raises the conditional delegated share by 2.0 percentage points, 13 percent of its 0.151 mean, with turnover showing the same pattern at 30 percent of the mean, while confirmation bias runs the other way, raising adoption and lowering the delegated share. Wealth operates on entry, raising adoption by 14.2 percentage points per standard deviation, and pre-adoption fund experience is the strongest deterrent to first adoption in the paper, carrying the largest first-stage coefficient in that design, -0.190 against a first-adoption rate of 7.7 percent. Adopters also differ from non-adopters on something we do not observe, and that unobserved difference raises the share they delegate rather than lowering it. The estimate of it turns negative when the sample is reweighted to platform proportions, so we offer the reading of it as trust as an interpretation rather than a finding. These patterns offer one account of why the adoption literature disagrees about who uses robo-advice (Baulkaran & Jain, 2023; Brenner & Meyll, 2020; Piehlmaier, 2022), since different samples recover different mixtures of entry, commitment, and lifecycle stages from a single use indicator.

Three mechanisms are consistent with why the margins separate, each tested on evidence the cross-section does not contain. Selective awareness accounts for the headline pattern: the mistakes the algorithm visibly corrects recede after adoption, most among initially high-bias and more engaged investors, while the biases delegation does not relax depress both margins and do not recede. Decision authority distinguishes robo-advising from automation as such, since within the same investor and month high-turnover and concentrated investors tilt toward the automatic plan and away from full delegation while engaged and attention-constrained investors are less likely to use AUTO and more likely to use RA. Reclaiming control makes delegation dynamic, with delegated shares falling

fastest for experienced adopters as the self-directed portfolio rebuilds. Recent performance enters the lifecycle margins in a structured way but absorbs none of the characteristic gradients.

The policy content follows from the mechanisms, not the magnitudes, and it is a prescription to test rather than a finding. If awareness gates entry, a platform seeking underdiversified or high-turnover investors should make those mistakes salient before adoption, and that is the intervention worth running against advertising algorithmic performance. If delegation is partly reversible, assets retained is an incomplete measure of whether a corrective product has worked.

Two limitations bound what we claim. The setting is one platform in one country. And the mechanism tests are observational: they support the readings of selective awareness, authority, trust, and learning without uniquely identifying them, and experiments varying the salience of own-portfolio mistakes or the decision rights transferred could do so directly. What may travel is the margin decomposition. Where participation requires delegation, entry and commitment can in principle respond to different characteristics, and target-date funds, managed accounts, and advisor-sold products share enough of that structure to be worth examining the same way. On the evidence here, the demand for financial advice is not well summarized by adoption alone.

References

Aghion, P., & Tirole, J. (1997). Formal and real authority in organizations. *Journal of Political Economy*, *105*(1), 1–29.

Agnew, J., Balduzzi, P., & Sundén, A. (2003). Portfolio choice and trading in a large 401(k) plan. *American Economic Review*, *93*(1), 193–215.

Ai, C., & Norton, E. C. (2003). Interaction terms in logit and probit models. *Economics Letters*, *80*(1), 123–129.

Arrow, K. J. (1962). The economic implications of learning by doing. *Review of Economic Studies*, *29*(3), 155–173.

Bailey, W., Kumar, A., & Ng, D. (2011). Behavioral biases of mutual fund investors. *Journal of Financial Economics*, *102*(1), 1–27.

Barber, B. M., & Odean, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance*, *55*(2), 773–806.

Barber, B. M., & Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *Quarterly Journal of Economics*, *116*(1), 261–292.

Barber, B. M., & Odean, T. (2002). Online investors: Do the slow die first? *Review of Financial Studies*, *15*(2), 455–488.

Barber, B. M., & Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies*, *21*(2), 785–818.

Barberis, N., Greenwood, R., Jin, L., & Shleifer, A. (2015). X-CAPM: An extrapolative capital asset pricing model. *Journal of Financial Economics*, *115*(1), 1–24.

Bartling, B., Fehr, E., & Herz, H. (2014). The intrinsic value of decision rights. *Econometrica*, *82*(6), 2005–2039.

Baulkaran, V., & Jain, P. (2023). Who uses robo-advising and how? *Financial Review*, *58*(1), 65–89.

Belanche, D., Casaló, L. V., & Flavián, C. (2019). Artificial intelligence in fintech: Understanding robo-advisors adoption among customers. *Industrial Management & Data Systems*, *119*(7), 1411–1430.

Benartzi, S., & Thaler, R. H. (2001). Naive diversification strategies in defined contribution saving plans. *American Economic Review*, *91*(1), 79–98.

Ben-Porath, Y. (1967). The production of human capital and the life cycle of earnings. *Journal of Political Economy*, *75*(4), 352–365.

Bhattacharya, U., Hackethal, A., Kaesler, S., Loos, B., & Meyer, S. (2012). Is unbiased financial advice to retail investors sufficient? Answers from a large field study. *Review of Financial Studies*, *25*(4), 975–1032.

Bianchi, M., & Brière, M. (2026). Human-robot interactions in investment decisions. *Management Science*, *72*(1), 14–31.

Bloomberg News. (2021, March). *Ant, Vanguard-backed robo adviser hits 1 million users in China*. Retrieved July 29, 2026, from https://www.bloomberg.com/news/articles/2021-03-18/ant-vanguard-backed-robo-adviser-hits-1-million-users-in-china

Blume, M. E., & Friend, I. (1975). The asset structure of individual portfolios and some implications for utility functions. *Journal of Finance*, *30*(2), 585–603.

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

Calvet, L. E., Campbell, J. Y., & Sodini, P. (2007). Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy*, *115*(5), 707–747.

Calvet, L. E., Campbell, J. Y., & Sodini, P. (2009). Fight or flight? Portfolio rebalancing by individual investors. *Quarterly Journal of Economics*, *124*(1), 301–348.

Campbell, J. Y. (2006). Household finance. *Journal of Finance*, *61*(4), 1553–1604.

Chalmers, J., & Reuter, J. (2020). Is conflicted investment advice better than no advice? *Journal of Financial Economics*, *138*(2), 366–387.

Chang, T. Y., Solomon, D. H., & Westerfield, M. M. (2016). Looking for someone to blame: Delegation, cognitive dissonance, and the disposition effect. *Journal of Finance*, *71*(1), 267–302.

Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2009). Reinforcement learning and savings behavior. *Journal of Finance*, *64*(6), 2515–2534.

D'Acunto, F., Prabhala, N., & Rossi, A. G. (2019). The promises and pitfalls of robo-advising. *Review of Financial Studies*, *32*(5), 1983–2020.

D'Acunto, F., & Rossi, A. G. (2023). Robo-advice: Transforming households into rational economic agents. *Annual Review of Financial Economics*, *15*, 543–563.

D'Hondt, C., De Winne, R., Ghysels, E., & Raymond, S. (2020). Artificial intelligence alter egos: Who might benefit from robo-investing? *Journal of Empirical Finance*, *59*, 278–299.

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

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, *64*(3), 1155–1170.

Egan, M., Matvos, G., & Seru, A. (2019). The market for financial adviser misconduct. *Journal of Political Economy*, *127*(1), 233–295.

Fama, E. F., & MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, *81*(3), 607–636.

Fehr, E., Herz, H., & Wilkening, T. (2013). The lure of authority: Motivation and incentive effects of power. *American Economic Review*, *103*(4), 1325–1359.

Foerster, S., Linnainmaa, J. T., Melzer, B. T., & Previtero, A. (2017). Retail financial advice: Does one size fit all? *Journal of Finance*, *72*(4), 1441–1482.

Gennaioli, N., Shleifer, A., & Vishny, R. (2015). Money doctors. *Journal of Finance*, *70*(1), 91–114.

Gervais, S., & Odean, T. (2001). Learning to be overconfident. *Review of Financial Studies*, *14*(1), 1–27.

Glaser, M., & Weber, M. (2007). Why inexperienced investors do not learn: They do not know their past portfolio performance. *Finance Research Letters*, *4*(4), 203–216.

Goetzmann, W. N., & Kumar, A. (2008). Equity portfolio diversification. *Review of Finance*, *12*(3), 433–463.

Greenwood, R., & Shleifer, A. (2014). Expectations of returns and expected returns. *Review of Financial Studies*, *27*(3), 714–746.

Greig, F., Ramadorai, T., Rossi, A. G., Utkus, S. P., & Walther, A. (2025). Human financial advice in the age of automation [Working paper, SSRN 4301514].

Grinblatt, M., & Keloharju, M. (2009). Sensation seeking, overconfidence, and trading activity. *Journal of Finance*, *64*(2), 549–578.

Guiso, L., Sapienza, P., & Zingales, L. (2008). Trusting the stock market. *Journal of Finance*, *63*(6), 2557–2600.

Hackethal, A., Haliassos, M., & Jappelli, T. (2012). Financial advisors: A case of babysitters? *Journal of Banking & Finance*, *36*(2), 509–524.

Han, B., & Kumar, A. (2013). Speculative retail trading and asset prices. *Journal of Financial and Quantitative Analysis*, *48*(2), 377–404.

Hao, R., Hu, C., Xu, X., & Zhang, Y. (2022). Does robo-advising improve retail investors' financial sophistication? [Working paper, SSRN 4230191; circulated earlier as "Beyond Performance: The Financial Education Role of Robo-Advising"].

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, *47*(1), 153–161.

Holzmeister, F., Holmén, M., Kirchler, M., Stefan, M., & Wengström, E. (2023). Delegation decisions in finance. *Management Science*, *69*(8), 4828–4844.

Hong, C. Y., Lu, X., & Pan, J. (2025). Fintech platforms and mutual fund distribution. *Management Science*, *71*(1), 488–517.

Huberman, G. (2001). Familiarity breeds investment. *Review of Financial Studies*, *14*(3), 659–680.

Kim, H. H., Maurer, R., & Mitchell, O. S. (2016). Time is money: Rational life cycle inertia and the delegation of investment management. *Journal of Financial Economics*, *121*(2), 427–447.

Klemperer, P. (1995). Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *The Review of Economic Studies*, *62*(4), 515–539.

Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, *77*(6), 1121–1134.

Kumar, A. (2009). Who gambles in the stock market? *Journal of Finance*, *64*(4), 1889–1933.

Lakonishok, J., Shleifer, A., & Vishny, R. W. (1992). The impact of institutional trading on stock prices. *Journal of Financial Economics*, *32*(1), 23–43.

Li, K., & Prabhala, N. R. (2007). Self-selection models in corporate finance. In B. E. Eckbo (Ed.), *Handbook of empirical corporate finance* (Vol. 1). North-Holland.

Linnainmaa, J. T. (2011). Why do (some) households trade so much? *Review of Financial Studies*, *24*(5), 1630–1666.

Linnainmaa, J. T., Melzer, B. T., & Previtero, A. (2021). The misguided beliefs of financial advisors. *Journal of Finance*, *76*(2), 587–621.

Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, *151*, 90–103.

Manski, C. F., & Lerman, S. R. (1977). The estimation of choice probabilities from choice based samples. *Econometrica*, *45*(8), 1977–1988.

Mitchell, O. S., & Utkus, S. P. (2022). Target-date funds and portfolio choice in 401(k) plans. *Journal of Pension Economics and Finance*, *21*(4), 519–536.

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, *55*(3), 703–708.

Nicolosi, G., Peng, L., & Zhu, N. (2009). Do individual investors learn from their trading experience? *Journal of Financial Markets*, *12*(2), 317–336.

Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance*, *53*(5), 1775–1798.

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, *39*(2), 230–253.

Philippon, T. (2019). *On fintech and financial inclusion* (Working Paper No. 26330). National Bureau of Economic Research.

Piehlmaier, D. M. (2022). Overconfidence and the adoption of robo-advice: Why overconfident investors drive the expansion of automated financial advice. *Financial Innovation*, *8*(1), 14.

Polkovnichenko, V. (2005). Household portfolio diversification: A case for rank-dependent preferences. *Review of Financial Studies*, *18*(4), 1467–1502.

Reher, M., & Sokolinski, S. (2024). Robo advisors and access to wealth management. *Journal of Financial Economics*, *155*, 103829.

Rossi, A. G., & Utkus, S. (2024). The diversification and welfare effects of robo-advising. *Journal of Financial Economics*, *157*, 103869.

Roth, J., Sant'Anna, P. H. C., Bilinski, A., & Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, *235*(2), 2218–2244.

Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty*, *1*(1), 7–59.

Scherer, B., & Lehner, S. (2023). Trust me, i am a robo-advisor. *Journal of Asset Management*, *24*(2), 85–96.

Scherer, B., & Lehner, S. (2025). What drives robo-advice? *Journal of Empirical Finance*, *80*, 101574.

Semykina, A., & Wooldridge, J. M. (2010). Estimating panel data models in the presence of endogeneity and selection. *Journal of Econometrics*, *157*(2), 375–380.

Seru, A., Shumway, T., & Stoffman, N. (2010). Learning by trading. *Review of Financial Studies*, *23*(2), 705–739.

Shefrin, H., & Statman, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance*, *40*(3), 777–790.

Sialm, C., Starks, L. T., & Zhang, H. (2015). Defined contribution pension plans: Sticky or discerning money? *Journal of Finance*, *70*(2), 805–838.

Sicherman, N., Loewenstein, G., Seppi, D. J., & Utkus, S. P. (2016). Financial attention. *Review of Financial Studies*, *29*(4), 863–897.

Statman, M., Thorley, S., & Vorkink, K. (2006). Investor overconfidence and trading volume. *Review of Financial Studies*, *19*(4), 1531–1565.

Van Nieuwerburgh, S., & Veldkamp, L. (2010). Information acquisition and under-diversification. *Review of Economic Studies*, *77*(2), 779–805.

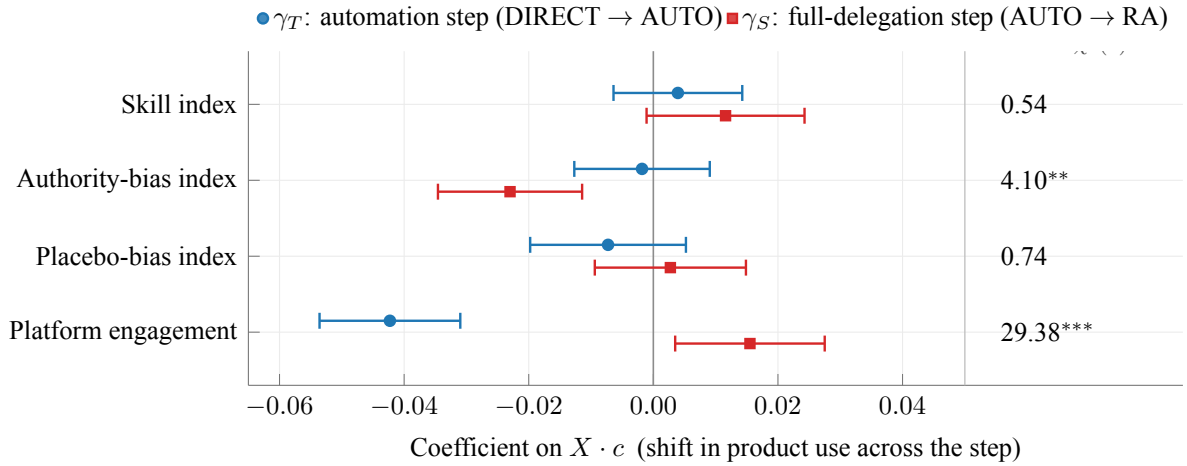
Vella, F. (1998). Estimating models with sample selection bias: A survey. *The Journal of Human Resources*, *33*(1), 127–169.

von Gaudecker, H.-M. (2015). How does household portfolio diversification vary with financial literacy and financial advice? *Journal of Finance*, *70*(2), 489–507.

Wooldridge, J. M. (1995). Selection corrections for panel data models under conditional mean independence assumptions. *Journal of Econometrics*, *68*(1), 115–132.

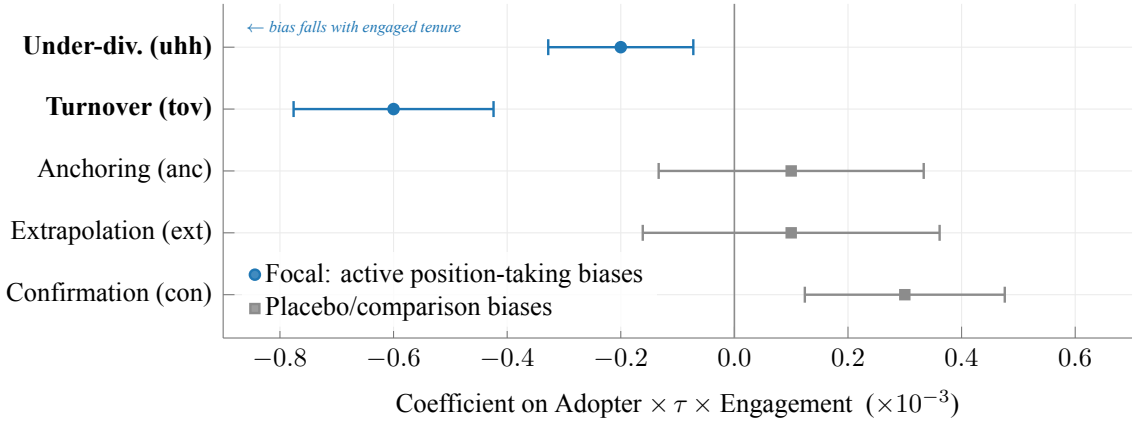
Figures

Figure 1: Decision-Authority Mechanism: What Moves Investors Across Each Delegation Step



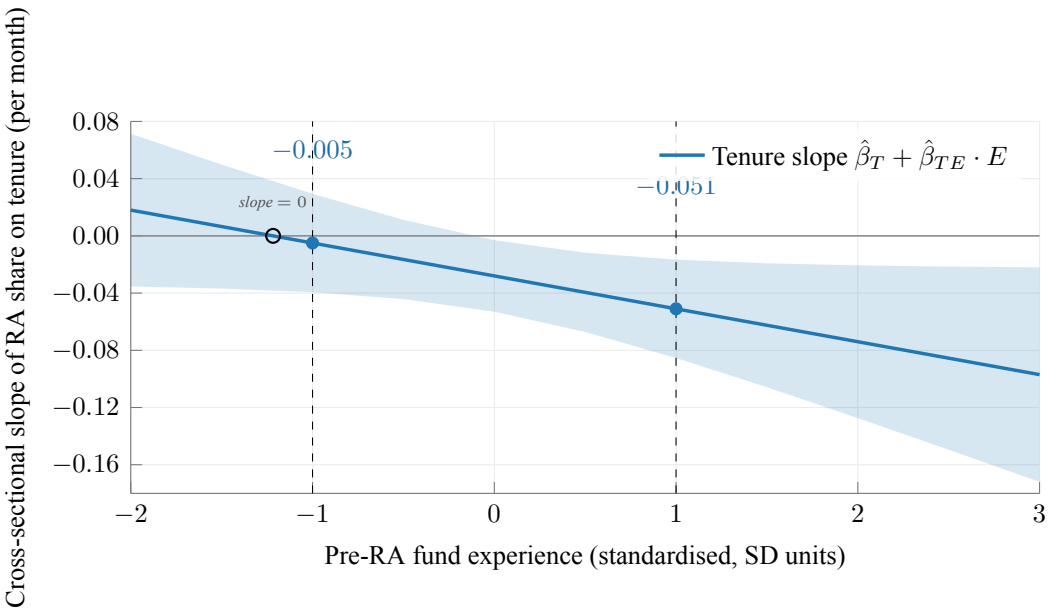
Notes. Each row reports two coefficients from the headline stacked linear-probability specification using the six-month pre-use window and the identifying-only sample. The model is estimated on the investor-month × product panel with investor-month choice-set fixed effects. γ_T (blue circles) is the coefficient on $X \cdot c_T$ and measures the difference in the association of X with AUTO versus DIRECT use. γ_S (red squares) is the coefficient on $X \cdot c_S$ and measures the incremental difference for RA versus AUTO use. Horizontal bars are 95% confidence intervals computed from the reported two-way clustered standard errors. The right-hand column reports the Wald $\chi^2(1)$ test of $\gamma_T = \gamma_S$ for each index (** $p < 0.05$, *** $p < 0.01$). Overlap of the two confidence intervals is not the equality test, because the Wald test uses the joint covariance of the two coefficients.
Reading the figure. The skill index is not associated with the automation step ($\gamma_T = 0.004$, $t = 0.75$), while its coefficient for the full-delegation step is positive and marginally significant ($\gamma_S = 0.012$, $t = 1.79$); the two coefficients are not statistically different (Wald $\chi^2 = 0.54$, $p = 0.461$). The authority-bias index is approximately zero for the automation step but negative for the full-delegation step ($\gamma_S = -0.023$, $t = -3.90$), and the two step coefficients differ at the 5% level (Wald $\chi^2 = 4.10$, $p = 0.043$). The placebo-bias index is statistically indistinguishable from zero at both steps and shows no cross-step difference. Platform engagement exhibits the sharpest split: it is strongly negative for the automation step ($\gamma_T = -0.042$, $t = -7.34$) and positive for the full-delegation step ($\gamma_S = 0.016$, $t = 2.53$), with a highly significant equality rejection (Wald $\chi^2 = 29.38$, $p < 0.001$). These patterns are descriptive associations consistent with the decision-authority mechanism, not causal effects.

Figure 2: Selective Awareness: Engagement-Dosed Bias Response to RA Tenure



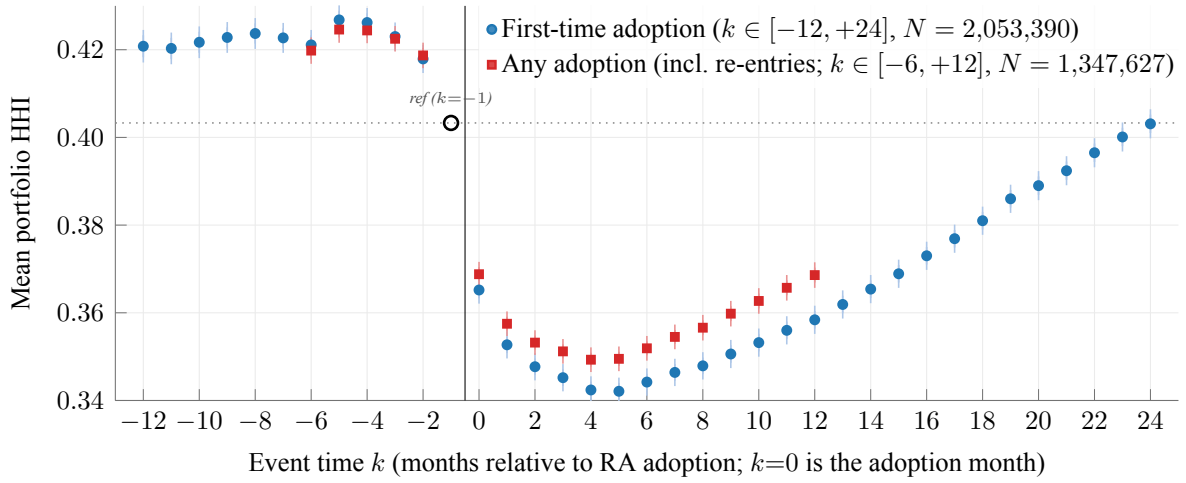
Notes. Each point is the triple-interaction coefficient Adopter $\times \tau \times$ Engagement (multiplied by 10^3) from the difference-in-differences specification with matched never-adopters in Panel A of Table 5; τ is months since first robo-advisor adoption (capped at 24) and Engagement is the standardised lagged 12-month login count. The coefficient measures how the bias trajectory over RA tenure steepens with platform engagement. Horizontal bars are 95% confidence intervals; standard errors are recovered as $|\text{coef}/t|$ from t-statistics two-way clustered by investor and calendar month. *Reading the figure.* For engaged adopters, the two focal biases that require active position-taking—under-diversification (-0.0002, $t = -3.07$) and turnover (-0.0006, $t = -6.68$)—decline significantly faster with RA tenure. Among the comparison biases, anchoring and extrapolation are statistically indistinguishable from zero, consistent with the engagement dose not producing a blanket bias improvement. Confirmation bias is the one exception: it carries a small, significantly *positive* coefficient ($+0.0003$, $t = 3.34$), i.e. it does not decline. The contrast—focal biases falling, comparison biases flat or rising—is consistent with selective awareness for the biases tied to active position-formation rather than across-the-board improvement.

Figure 3: Reclaiming Control: Tenure Slope of RA Share by Pre-RA Experience



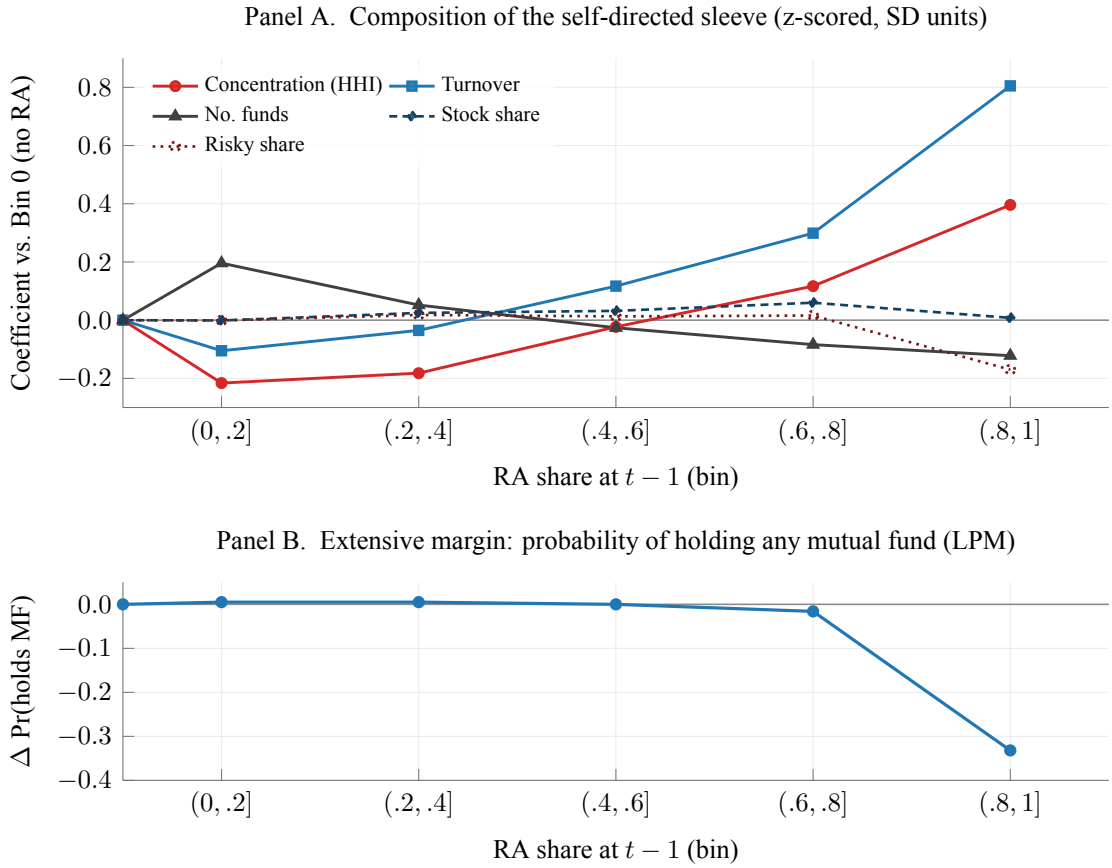
Notes. The figure plots the cross-sectional slope of RA share (RA assets / (RA + MF), among RA-active adopter-months) on robo-advisor tenure, as a function of standardised pre-RA fund experience E. Because estimation is Fama–MacBeth (monthly cross-sectional regressions), this is the average within-month gradient of RA share with respect to tenure—the RA-share difference associated with a one-month difference in tenure across adopters—rather than a within-investor time derivative, and it is a local slope that should not be extrapolated over many months. Coefficients are from column (1) of Table 8 (linear interaction, Fama–MacBeth OLS): $\hat{\beta}_T = -0.028$ (Tenure main effect, $t = -2.19$) and $\hat{\beta}_{TE} = -0.023$ (Pre-RA experience × Tenure, $t = -1.91$); the plotted quantity is $\hat{\beta}_T + \hat{\beta}_{TE} \cdot E$. The shaded band is a 95% confidence interval from the delta method, treating $\hat{\beta}_T$ and $\hat{\beta}_{TE}$ as approximately independent and recovering standard errors as $|\hat{\beta}/t|$. *Reading the figure.* The tenure slope is essentially zero for the least-experienced adopters ($E = -1$ SD: -0.005, within the confidence band of zero) and turns increasingly negative as pre-RA experience rises ($E = +1$ SD: -0.051). The slope crosses zero at $E \approx -1.2$ SD, so for all but the least-experienced adopters longer RA tenure is associated with a lower RA share of wallet, and the gradient is most negative for the most fund-experienced investors. This is the reclaiming-control pattern: investors with more pre-RA fund-investing experience exhibit lower delegated shares at longer RA tenures. The experience×tenure interaction is itself significant only at the 10% level ($\hat{\beta}_{TE} = -0.023$, $t = -1.91$), so the *steepening* of the decline with experience is suggestive; the tenure decline at and above mean experience is significant at the 5–1% level (e.g. $t = -2.9$ at $E = +1$). We plot the slope rather than fitted RA-share levels because the linear specification produces predicted shares outside $[0, 1]$ at extreme combinations of tenure and experience; the slope plot is invariant to that boundary issue.

Figure 4: What Changes When Investors Adopt RA?



Notes. Mean portfolio Herfindahl index (HHI) by event time relative to RA adoption. Each level is computed as $\bar{y}_{k=-1} + \hat{\beta}_k$, where $\hat{\beta}_k$ is the event-study coefficient from a within-investor fixed-effects regression of HHI on event-time dummies and $\bar{y}_{k=-1}$ is the cross-sectional mean of HHI in the month before adoption (the omitted reference category in the regression: 0.4033 for first-time adoption, 0.4066 for any adoption). Vertical bars are 95% confidence intervals, computed as $1.96 \times \text{se}(\hat{\beta}_k)$ and treating the baseline as a known constant. The vertical dashed line marks the adoption transition between $k = -1$ and $k = 0$; the horizontal dotted line marks the first-time-adoption baseline of 0.4033. The open marker at $k = -1$ is the omitted reference category. *Two series.* The blue circles plot the *first-time adoption* specification, anchoring each investor at her first-ever RA adoption ($k \in [-12, +24]$, N = 2,053,390). The red squares plot the *any-adoption* specification that stacks every adoption event including re-entries ($k \in [-6, +12]$, N = 1,347,627). The two series have different sample windows because the first-time specification can look further back from $k = 0$ (more pre-adoption history is available per investor), while the any-adoption specification has shorter usable pre/post windows once re-entry events are stacked. *Reading.* Mean HHI rises modestly in the year before first adoption, peaking around 0.427 at $k = -5$, then drops sharply to 0.365 in the adoption month, troughs at 0.342 four to five months after adoption, and drifts back toward the pre-adoption baseline thereafter, reaching 0.403 at $k = +24$ ($\hat{\beta}_{24} = -0.0002$, $|t| < 1$, indistinguishable from the reference). The any-adoption series tracks the first-time series closely over its shorter window, suggesting re-entries reproduce the same diversification pattern. Companion outcomes (RA share and risky RA share) are reported in the main event-study table.

Figure 5: Robo-Advisor Dose-Response on the Self-Directed Mutual-Fund Sleeve



Notes. Each marker is the within-investor coefficient on a lag-1 RA-share bin from the two-way fixed-effects regressions in Table 9, relative to the omitted no-RA reference (Bin 0, $rsh_{t-1} = 0$). Points are plotted at the bin midpoint, and each series is anchored at the origin, where Bin 0 (the omitted no-RA reference) has a coefficient of zero by construction. Connecting lines join the five estimated bin coefficients and serve only as visual guides; the bins are categorical, not a continuous dose. *Panel A* shows the composition of the self-directed mutual-fund sleeve: outcomes are winsorised at 1/99 and z-scored within month, so coefficients are in within-month standard-deviation units. As RA intensity rises, the residual self-directed sleeve becomes more concentrated (HHI), higher-turnover, and held in fewer funds; stock share rises modestly and risky share falls only at the highest dose. *Panel B* shows the extensive margin: the linear-probability coefficient for holding any mutual fund is flat across low and middle doses and collapses at the top bin (-0.332), i.e. investors with near-total RA share largely exit self-directed mutual funds altogether. Among columns (2)–(6), every bin coefficient is significant at the 1% level except: Stock share Bins 1 and 5 (insignificant); Risky share Bins 1, 3, and 4 (insignificant) and Bin 2 (significant at the 10% level only). In column (1), Bin 3 is insignificant. Standard errors are two-way clustered on investor and month. These are within-investor co-movements across the RA-share dose, not causal effects of RA usage.

Tables

Table 1: Summary Statistics for Variables Used in the Main Heckman and Robustness Analyses

Variable	N	Mean	SD	Min	p25	Median	p75	Max
Panel A. Demographics (Investor-Level)								
Age (years)	100,000	32.499	9.607	19.000	25.000	30.000	37.000	64.000
Risk Preference (0-5)	99,992	2.781	1.069	0.000	2.000	3.000	4.000	5.000
Female dummy	100,000	0.425	—	0.000	—	—	—	1.000
Panel B. Assets and Allocation (Investor-Month)								
Total Financial Assets (CNY)	4,900,000	95,245	174,510	0.000	6,534.730	30,220	99,531	1.37M
Robo-Advisor Assets (CNY)	4,900,000	2,194.960	8,526.560	0.000	0.000	0.000	199.460	102,820
Total Mutual-Fund Assets (CNY)	4,900,000	66,238	134,884	0.000	1,833.940	15,482	63,181	1.11M
Baby/Money-Market Assets (CNY)	4,900,000	20,630	48,786	0.000	167.667	2,656.380	16,043	393,365
Log(Total Financial Assets)	4,900,000	9.824	2.571	0.000	8.785	10.316	11.508	14.132
RA Share of TFA	4,837,721	0.030	0.087	0.000	0.000	0.000	0.007	0.712
RA Share of (RA + MF), unconditional	4,492,147	0.053	0.148	0.000	0.000	0.000	0.017	1.000
RA Share of (RA + MF), | RA active	1,571,475	**0.151**	0.219	0.000	0.013	0.057	0.185	1.000
Risky RA Share, unconditional	4,452,494	0.079	0.189	0.000	0.000	0.000	0.038	1.000
Risky RA Share, | risky RA active	—	0.225	0.262	—	0.028	0.115	0.331	—
Login Count (last 12 months)	4,899,785	1,758	1,402	37.000	739.000	1,382	2,378	7,977
Months Since Registration	4,900,000	112.691	42.284	11.000	82.000	109.000	141.000	229.000
Panel C. Behavioral Biases, Diversification, and Trading Skills								
Behavioral biases (winsorised 1%/99% within month; natural units, not standardized)								
Anchoring Bias	4,670,815	0.491	0.428	0.000	0.000	0.500	1.000	1.000
Total Turnover	4,402,068	0.291	0.390	3.20e-05	0.021	0.096	0.420	4.610
Attention Bias	4,670,815	0.416	0.412	0.000	0.000	0.323	0.913	1.000
Disposition Effect	4,367,178	0.022	0.376	-1.000	0.000	0.000	0.000	1.000
Underdiversification (HHI)	4,717,162	0.423	0.318	0.032	0.162	0.319	0.604	1.000
Extrapolation Bias	4,670,815	0.187	0.325	0.000	0.000	0.000	0.247	1.000
Gambling Bias	4,670,815	0.125	0.267	0.000	0.000	0.000	0.076	1.000
Confirmation Bias	3,619,997	0.821	0.340	0.000	0.900	0.990	1.000	1.000
Herding Bias	4,670,815	0.871	0.284	0.000	0.950	1.000	1.000	1.000
Portfolio diversification								
Number of Fund Holdings	4,717,162	10.493	13.982	1.000	3.000	6.000	12.000	94.000
Effective N ($1/\sum w^2$)	4,717,162	4.929	5.157	1.000	1.655	3.139	6.163	31.742
Weighted Portfolio Correlation	3,568,973	0.367	0.273	-0.329	0.158	0.394	0.572	0.942
Skills and experience								
Fund Experience (months)	4,717,162	23.215	13.420	1.000	12.000	23.000	34.000	49.000
Selection Skill	4,685,235	-0.002	0.024	-0.233	-0.011	-0.001	0.007	0.138
Timing Rate	4,571,144	2.14e-04	0.015	-0.288	-0.002	1.75e-04	0.002	0.312
RA Experience (months)	2,377,855	14.426	10.932	0.000	6.000	11.000	21.000	49.000
Panel D. Returns (Winsorised at 1%/99% Within Month; Natural Units, Not Standardized)								
Self-directed fund returns								
Fund Raw Return	4,243,288	-0.003	0.055	-0.500	-0.029	-3.79e-04	0.019	0.305
Fund Excess Return vs CSI300	4,233,456	4.46e-04	0.080	-0.524	-0.038	0.006	0.045	0.337
Fund β-Adj Return vs CSI300	3,308,592	-0.004	0.052	-0.223	-0.027	-0.002	0.012	0.298
Fund Category-Adjusted Return	4,180,163	-0.003	0.026	-0.245	-0.013	-0.002	0.007	0.143
Robo-advisor returns								
RA Raw Return	1,492,840	-0.002	0.031	-0.287	-0.014	0.000	0.006	0.165
RA Excess Return vs CSI300	1,492,840	0.002	0.065	-0.311	-0.025	0.007	0.039	0.197
RA β-Adj Return vs CSI300	1,172,397	-0.002	0.031	-0.128	-0.014	-7.38e-04	0.007	0.160
RA Strategy-Adjusted Return	1,485,437	-5.83e-04	—	—	—	—	—	—
Panel E. Auto-Invest Usage and Product Overlap								
Auto-Invest Usage (monthly)	4,900,000	0.329	—	0.000	—	—	—	1.000
Auto-Invest Share (if active)	1,614,024	0.733	0.385	2.00e-06	0.353	1.000	1.000	1.000
RA × auto-invest cross-tabulation of the 4,900,000 investor-months:								
Both RA and auto-invest	683,373	13.95%	—	—	—	—	—	—
RA only	888,102	18.12%	—	—	—	—	—	—
Auto-invest only	930,651	18.99%	—	—	—	—	—	—

Continued on next page

Table 1 (continued)

Variable	N	Mean	SD	Min	p25	Median	p75	Max
Neither	2,397,874	48.94%	—	—	—	—	—	—
Panel F. Adoption Lifecycle (Investor-Month Transitions)								
Month-to-month state transitions across 4,900,000 investor-months:								
Stay out (non-RA → non-RA)	3,287,884	67.10%	—	—	—	—	—	—
Continue (RA → RA)	1,491,648	30.44%	—	—	—	—	—	—
Enter (non-RA → RA)	79,827	1.63%	—	—	—	—	—	—
Exit (RA → non-RA)	40,641	0.83%	—	—	—	—	—	—
Conditional rates:								
First-adoption rate (per at-risk investor-month)	—	1.43%	—	—	—	—	—	—
Continuation rate (| RA at $t-1$)	1,532,289	97.35%	—	—	—	—	—	—
Re-entry rate (| RA at $\leq t-2$ only)	752,774	1.30%	—	—	—	—	—	—
RA-Active Months per Adopter	70,034	22.439	14.172	1.000	10.000	18.000	36.000	49.000

Notes. Summary moments for all variables used in the main Heckman regressions and robustness tests, computed on the investor–month panel covering 100,000 investors over 49 analysis months (4.9 million investor-months total); the data span September 2020 – September 2024, with the first month reserved for constructing one-month-lagged regressors. **All continuous variables are winsorised at the 1st and 99th percentiles within each calendar month before the moments are computed, but they are *not* standardized: all statistics are in natural units in the table.** The within-month z-scoring used in estimation is applied inside the regression pipeline only, after these moments. Because winsorisation caps vary month by month, pooled minima and maxima reflect the most extreme monthly caps and can lie outside the pooled 1st–99th percentile range. **Panel A** reports time-invariant investor-level demographics computed once per investor. **Panel B** reports investor–month asset balances in Chinese yuan (CNY) and allocation shares. The conditional row *RA Share of (RA + MF) | RA active* (mean 0.151, N = 1,571,475) is the dependent variable in the outcome stage of the Heckman model and the RA Share regressions; the unconditional version above it (which includes all RA-inactive observations as zeros) has mean 0.053. The Risky RA Share, the dependent variable in the risky-allocation robustness specifications, excludes low-risk RA strategies from the numerator and money-market funds from the denominator; its conditional row restricts to months with positive risky RA holdings. Login Count (last 12 months) is reported here for the engagement specifications; the median investor logs in roughly ∼1,382 times per year and the inter-quartile range spans 739–2,378 logins. **Panel C** reports behavioral biases, portfolio-diversification measures, and trading skills; biases take natural values on their constructed ranges (most on $[0, 1]$, Disposition Effect on $[-1, 1]$); Effective N and Weighted Portfolio Correlation are computed on investors with at least two fund holdings. **Panel D** reports fund and robo-advisor monthly returns: (i) raw, (ii) excess vs the CSI300 index, (iii) β-adjusted excess (using a 36-month rolling market β with a 12-month minimum), and (iv) category-adjusted (self-directed funds) or strategy-peer-adjusted (RA). Only N and the mean were reported for the RA Strategy-Adjusted Return in the source output. **Panel E** reports auto-invest usage and product overlap: the monthly investor-level participation rate (0.329) and, conditional on participation, the share of fund purchases executed via automated plans (0.733). The cross-tabulation partitions all 4,900,000 investor-months by contemporaneous RA and auto-invest status; the four cells (Both 683,373; RA only 888,102; Auto-invest only 930,651; Neither 2,397,874) are mutually exclusive and sum exactly to the panel. **Panel F** reports adoption-lifecycle transitions across consecutive investor-months. The four mutually exclusive month-to-month transitions sum to the full panel of 4,900,000 rows. Conditional rates convert these counts into hazard-style rates: the continuation rate conditions on the 1,532,289 at-risk months with RA holdings in $t-1$ (of which 1,491,648 continue and 40,641 exit); the re-entry rate conditions on the 752,774 at-risk months of prior adopters inactive in $t-1$; the first-adoption rate is per at-risk investor-month among never-yet-adopters. 70.0% of investors adopt RA at some point during the sample window (70,034 of 100,000), and conditional on ever adopting, investors hold RA for 22.4 months on average. Sample sizes differ across rows because (i) RA and auto-invest moments are conditional on the relevant product being active; (ii) return moments require non-missing NAV data; (iii) bias and skill moments require minimum fund-holding histories. N is reported as the total number of non-missing observations at the relevant unit (investor or investor-month). All statistics are unweighted.

Table 2: Main Heckman Results

	Panel 1: Baseline (T=48)		Panel 2: +Skills (T=47)		Panel 3: +Biases (T=36)	
Variable	Probit ME	Outcome (RA_{share})	Probit ME	Outcome (RA_{share})	Probit ME	Outcome (RA_{share})
Constant	-0.197***	0.054	-0.190***	-0.036	-0.152***	-0.030
	(-6.752)	(0.824)	(-7.663)	(-0.690)	(-12.755)	(-0.662)
Log(Total Financial Assets)$_{t-1}$	0.142***	-0.079***	0.134***	-0.049***	0.079***	-0.010
	(12.560)	(-4.230)	(11.754)	(-5.280)	(6.152)	(-1.555)
Risk Preference	0.029***	0.008*	0.022***	0.020***	0.019***	0.012***
	(5.616)	(1.725)	(4.261)	(5.410)	(6.600)	(9.274)
Age	0.004	-0.005**	0.003	-0.010*	0.016***	-0.002
	(0.770)	(-2.010)	(0.612)	(-1.791)	(19.588)	(-0.868)
Female	-0.042***	-0.013***	-0.043***	-0.020***	-0.052***	-0.015***
	(-17.139)	(-2.851)	(-19.983)	(-4.461)	(-19.614)	(-3.624)
Combined RA Dummy$_{t-1}$	0.224***	0.036**	0.222***	0.051***	0.218***	0.067***
	(24.914)	(2.248)	(29.158)	(3.577)	(90.186)	(4.187)
Baby Assets$_{t-1}$	-0.037***	0.014***	-0.034***	0.009***	-0.021***	0.001
	(-18.856)	(3.497)	(-16.496)	(2.926)	(-9.788)	(0.309)
Fund Experience (Months)$_{t-1}$			0.035***	-0.028***		
			(9.546)	(-6.063)		
Selection Ability$_{t-1}$			0.004*	0.000		
			(1.886)	(-0.073)		
Timing Ability$_{t-1}$			0.001	-0.001		
			(0.968)	(-1.376)		
Anchoring Bias$_{t-1}$					0.021***	0.002
					(12.172)	(1.057)
Total Turnover$_{t-1}$					-0.030***	0.045***
					(-14.832)	(5.886)
Speculative Turnover$_{t-1}$					-0.002***	-0.019***
					(-3.365)	(-8.348)
Attention Bias$_{t-1}$					-0.009***	-0.005***
					(-7.133)	(-4.544)
Disposition Effect$_{t-1}$					0.003***	-0.002***
					(5.211)	(-3.971)
Underdiversification HHI$_{t-1}$					-0.065***	0.020***
					(-99.824)	(3.811)
Extrapolation Bias$_{t-1}$					-0.005***	-0.001
					(-2.878)	(-1.356)
Gambling Bias$_{t-1}$					-0.009***	-0.003***
					(-5.833)	(-2.977)
Confirmation Bias$_{t-1}$					0.030***	-0.022***
					(26.037)	(-5.566)
Herding Bias$_{t-1}$					-0.006***	-0.003***
					(-7.422)	(-2.762)
Months Since Registration$_{t-1}$	0.011***		0.008***		0.014***	
	(10.538)		(7.802)		(10.981)	
Inverse Mills Ratio (λ)		0.158***		0.203***		0.212***
		(3.480)		(4.494)		(4.324)
Province Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2 / R^2	0.082	0.133	0.078	0.146	0.090	0.229
N (total)	4,799,616	1,564,155	4,471,044	1,539,119	3,306,408	1,230,489

Notes. Estimates are from unweighted Heckman two-stage (Tobit-2) selection models estimated monthly and averaged following Fama and MacBeth (1973). Reported selection-stage values are *approximate marginal effects* on the probability of RA adoption, computed as $\text{ME}_k = \phi(\Phi^{-1}(\bar{P}_s)) \cdot \hat{\beta}_k$, where $\hat{\beta}_k$ is the Fama–MacBeth average Probit coefficient and $\bar{P}_s$ is the RA-active rate in each specification's estimation sample (RA-adopter months ÷ selection-equation investor-months): $\bar{P}_s = 0.326$ (Baseline), 0.344 (+Skills), 0.372 (+Biases), with scaling factors $\phi(\Phi^{-1}(\bar{P}_s)) = 0.360$, 0.368, and 0.378 respectively. Marginal effects are computed from the unrounded coefficients and scaling factors and then rounded for display. All estimates are reported to three decimals; Newey–West HAC t-statistics (max lag = 3), computed from the time series of monthly coefficients, are in parentheses; t-statistics, stars, and significance are invariant to the scaling. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed). Panel 1 reports the baseline specification. Panel 2 adds investor skills (fund experience, selection ability, timing ability). Panel 3 adds behavioral biases. The selection (Probit) dependent variable is $\text{RA}_{it} = 1$ if the investor holds robo-advisor assets; the outcome dependent variable is RA_{share} = RA assets / (RA + mutual fund assets), observed only for RA-adopter months. As a robustness check, Table A.1 reports the same specifications estimated with calibrated inverse-probability weights. Table A.2 reports the sensitivity of these estimates to alternative placements of the exclusion restriction.

Table 3: First Adoption of Robo-Advisor and Auto-Invest: Pre-Window Construction

	Robo-Advisor		Auto-Invest	
Variable	Probit ME	Outcome (RA_{share})	Probit ME	Outcome (AUTO_{share})
Panel A: Demographics				
log(TFA), pre-window avg	0.029***	-0.041***	0.042***	-0.080***
	(3.709)	(-6.146)	(5.860)	(-2.680)
Risk preference	0.007***	-0.005	0.008***	-0.005
	(6.824)	(-1.258)	(7.096)	(-0.450)
Age at adoption	-0.002	-0.006*	-0.001	0.001
	(-0.578)	(-1.825)	(-0.483)	(0.055)
Female	-0.007**	0.026***	0.004	0.053***
	(-2.065)	(3.686)	(0.837)	(2.956)
Combined-RA share	0.007***	-0.005	0.002	0.011
	(3.249)	(-1.256)	(0.838)	(0.686)
Baby assets	-0.013*	0.020***	-0.017***	0.006
	(-1.884)	(2.995)	(-5.280)	(0.401)
Panel B: Skills				
Fund experience	-0.190***	0.027	-0.137***	0.076
	(-6.196)	(0.562)	(-6.707)	(0.821)
Selection score (TM)	-0.013***	-0.013	-0.010**	0.044
	(-3.453)	(-0.790)	(-2.446)	(1.326)
Timing score (TM)	0.014**	-0.002	0.012	0.009
	(2.028)	(-0.407)	(1.442)	(0.462)
Panel C: Behavioural biases				
Anchoring	-0.030***	-0.032***	-0.031***	-0.027
	(-3.841)	(-2.620)	(-5.058)	(-1.112)
Turnover	-0.011*	0.005	-0.013***	-0.044**
	(-1.928)	(0.593)	(-4.030)	(-2.592)
Attention	-0.018***	0.006	-0.017***	-0.004
	(-3.575)	(1.109)	(-7.589)	(-0.212)
Disposition effect	0.002	-0.003	0.002	0.005
	(1.211)	(-1.161)	(1.282)	(0.582)
Under-diversification (HHI)	-0.005***	0.045***	-0.007***	0.051***
	(-3.437)	(8.034)	(-3.680)	(4.435)
Extrapolation	0.010	-0.005	-0.002	0.000
	(1.070)	(-1.455)	(-0.731)	(0.024)
Gambling	0.001	0.003	0.003	0.008
	(0.935)	(0.768)	(0.873)	(0.479)
Confirmation	0.030***	-0.040***	0.025***	0.025
	(4.801)	(-4.648)	(7.715)	(1.175)
Herding	-0.007*	-0.001	-0.009***	-0.013
	(-1.949)	(-0.590)	(-8.701)	(-0.731)
Months since reg., $t-1$	0.004		0.002	
	(0.722)		(1.500)	
Inverse Mills ratio		-0.081		-0.090
		(-0.989)		(-0.827)
Province Dummies	Yes	Yes	Yes	Yes
Pseudo R^2 / R^2	0.195	—	0.196	—
N	254,744	19,368	82,592	6,902

Notes. The table reports unweighted Heckman two-step (Tobit-2) estimates of first adoption of the robo-advisor (RA) and auto-invest (AUTO) products, estimated monthly and aggregated by Fama–MacBeth. The RA columns set $D_{i,t} = 1$ in the investor's first robo-advisor month with Stage 2 outcome $rsh_{i,t} = RA_{i,t}/(RA_{i,t} + MF_{i,t})$; the AUTO columns set $D_{i,t} = 1$ in the first auto-invest month with Stage 2 outcome $ash_{i,t} = auto_buy_{i,t}/total_buy_{i,t}$. Stage 2 in each event is estimated on the $D = 1$ cross-section. Both events use the same pre-window construction. Female is binary. Stage-1 (selection) values are approximate marginal effects on the first-adoption probability, $\text{ME}_k = \phi(\Phi^{-1}(\bar{P}))\hat{\beta}_k$, scaled by each event's own first-adoption rate in the pre-window estimation sample, – $\bar{P}_{\text{RA}} = 0.077$ (factor 0.144) and $\bar{P}_{\text{AUTO}} = 0.084$ (factor 0.154). Stage-2 outcome coefficients are unscaled; t-statistics and significance are invariant to the scaling. The reported specification adds investor skills (Panel B) and behavioural biases (Panel C) to the demographics block (Panel A); panels are organised by variable group for readability. Baseline-only and baseline + skills specifications, the lag-1 hazard-style design, and the fund-return specifications are reported in the Online Appendix. Newey–West HAC t-statistics (lag = $\min(3, T-1)$) in parentheses below estimates. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed). Province dummies included, not reported.

Table 4: Decision-Authority Mechanism: Stacked Linear Probability Model with Investor-Month Fixed Effects

Variable	Automation step DIRECT → AUTO (γ_T)	Full-delegation step AUTO → RA (γ_S)	Equality test $\chi^2(1)$
Specification (1). Composite indices and pre-use platform engagement			
Skill composite index	0.004	0.012*	0.542
	(0.747)	(1.792)	[0.461]
Authority-bias composite index	-0.002	-0.023***	4.102**
	(-0.327)	(-3.897)	[0.043]
Placebo-bias composite index	-0.007	0.003	0.735
	(-1.140)	(0.442)	[0.391]
Pre-use login frequency	-0.042***	0.016**	29.379***
	(-7.337)	(2.532)	[<0.001]
Within R^2	0.2672	0.2672	0.2672
Specification (2). Disaggregated authority-relevant skills and biases			
Selection ability	-0.006	0.014**	3.651*
	(-1.069)	(2.308)	[0.056]
Timing ability	0.009*	0.003	0.431
	(1.829)	(0.557)	[0.511]
Confirmation bias	0.036***	0.012	3.722*
	(5.449)	(1.640)	[0.054]
Total turnover	0.035***	-0.032***	31.797***
	(5.069)	(-4.913)	[<0.001]
Speculative turnover	-0.031***	-0.013**	3.539*
	(-5.832)	(-2.121)	[0.060]
Herding bias	-0.015**	-0.001	1.563
	(-2.392)	(-0.163)	[0.211]
Disposition effect	0.005	-0.009	1.662
	(0.947)	(-1.394)	[0.197]
Within R^2	0.2700	0.2700	0.2700
Specification (3). Disaggregated placebo biases			
Under-diversification (HHI)	0.058***	-0.033***	69.600***
	(9.101)	(-6.067)	[<0.001]
Gambling bias	0.000	-0.013**	1.114
	(-0.009)	(-2.012)	[0.291]
Anchoring bias	-0.038***	0.048***	57.794***
	(-5.979)	(8.165)	[<0.001]
Attention bias	-0.042***	0.026***	29.630***
	(-6.199)	(3.993)	[<0.001]
Extrapolation bias	0.010	-0.020***	6.096**
	(1.569)	(-3.066)	[0.014]
Within R^2	0.2708	0.2708	0.2708
Baseline-control interactions	Yes	Yes	Yes
Stacked N	672,228	672,228	672,228

Notes. The three products, ordered by the amount of decision authority the investor delegates to the platform, are DIRECT (self-directed mutual fund investing), AUTO (automatic-investment plans, in which the investor selects funds and the platform automates the timing of purchases), and RA (the robo-advisor, in which the algorithm selects funds and manages the portfolio). The dependent variable $use_{i,t,k}$ equals one if investor i holds product $k \in \{\text{DIRECT}, \text{AUTO}, \text{RA}\}$ in month t and zero otherwise; every investor-month is stacked into three rows, one per product alternative. Products may be held simultaneously, so this is a stacked multivariate linear probability model rather than a mutually exclusive multinomial-choice model. Two cumulative-authority indicators encode the stacking: c_T equals one for AUTO and RA and zero for DIRECT; c_S equals one for RA and zero otherwise. With entity fixed effects on the investor-month absorbing each investor-month's mean use rate, the coefficient γ_T on $X \cdot c_T$ identifies how covariate X shifts the investor across the DIRECT-to-AUTO step, and γ_S on $X \cdot c_S$ identifies the additional shift across the AUTO-to-RA step, where fund choice itself is surrendered). The Wald column tests $H_0 : \gamma_T = \gamma_S$ from the clustered covariance matrix. Type measures are per-investor averages over the six months immediately preceding the investor's first AUTO or RA use, requiring at least three non-missing pre-use months for every type variable. Baseline controls (Log TFA_{t-1}, risk preference, age, female indicator, baby-asset holdings_{t-1}), each interacted with c_T and c_S, are included in every specification but suppressed for brevity; those estimates, the product-step intercepts, and the underlying standard errors are reported in Table A.9. Estimation is with entity fixed effects on the investor-month; standard errors are two-way clustered on investor and month. The identifying-only sample drops investor-months in which the investor is on the same product across all three stacked rows. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed).

Table 5: Bias Trajectories over Robo-Advisor Tenure: Selective Learning

Panel A: Rolling 6-month trailing mean of bias (headline)

	Focal: predicted decline		Placebo: predicted null		
Regressor	Under-div. (uhh)	Turnover (tov)	Anchoring (anc)	Extrap. (ext)	Confirm. (con)
Within-adopter sample					
τ	-0.0013^{***}	-0.0029^{***}	$+0.0009^{***}$	$+0.0002$	-0.0005^{**}
	(-6.53)	(-11.23)	$(+3.16)$	$(+0.36)$	(-2.35)
$\tau\times$ High Under-div.	-0.0109^{***}	-0.0050^{***}	$+0.0014^{***}$	$+0.0012^{***}$	-0.0001
	(-20.22)	(-7.48)	$(+2.87)$	$(+3.15)$	(-0.24)
$\tau\times$ High Turnover	-0.0029^{***}	-0.0157^{***}	$+0.0010^{**}$	$+0.0004$	$+0.0017^{***}$
	(-7.77)	(-15.89)	$(+2.10)$	$(+1.26)$	$(+3.22)$
$\tau\times$ AOT	-0.0056^{***}	-0.0103^{***}	$+0.0008^{**}$	$+0.0005^{*}$	$+0.0010^{**}$
	(-16.66)	(-12.74)	$(+2.27)$	$(+1.81)$	$(+2.41)$
Difference-in-differences with matched never-adopters					
Adopter $\times\ \tau$	-0.0008^{***}	-0.0018^{***}	$+0.0015^{***}$	$+0.0011^{***}$	-0.0002
	(-6.20)	(-6.83)	$(+6.89)$	$(+3.81)$	(-1.20)
Adopter $\times\ \tau\ \times$ Engagement	-0.0002^{***}	-0.0006^{***}	$+0.0001$	$+0.0001$	$+0.0003^{***}$
	(-3.07)	(-6.68)	$(+0.84)$	$(+0.75)$	$(+3.34)$
Adopter $\times\ \tau\ \times$ High Under-div.	-0.0109^{***}	-0.0052^{***}	$+0.0013^{***}$	$+0.0011^{***}$	-0.0001
	(-20.49)	(-7.80)	$(+2.69)$	$(+2.90)$	(-0.15)
Adopter $\times\ \tau\ \times$ High Turnover	-0.0029^{***}	-0.0157^{***}	$+0.0010^{**}$	$+0.0004$	$+0.0017^{***}$
	(-7.69)	(-15.95)	$(+2.12)$	$(+1.25)$	$(+3.23)$
Adopter $\times\ \tau\ \times$ AOT	-0.0057^{***}	-0.0104^{***}	$+0.0008^{**}$	$+0.0005^{*}$	$+0.0011^{**}$
	(-16.70)	(-12.92)	$(+2.19)$	$(+1.67)$	$(+2.45)$
Observations (adopter-only)	1,321,291	1,321,291	1,321,291	1,321,291	1,150,704
Observations (DiD sample)	1,847,939	1,847,939	1,847,939	1,847,939	1,611,731

Notes for Panel A. The table reports two-way fixed-effects panel regressions of an investor-month behavioural bias measure on a tenure variable and its interactions. The dependent variable is the rolling six-month trailing mean of the bias measure ending at month $t-1$, so the month-t realisation is constructed entirely from history. Every regression includes investor and calendar-month fixed effects; standard errors are two-way clustered by investor and calendar month. The two focal biases, under-diversification (uhh) and turnover (tov), proxy for active opinion-taking and are predicted to decline as robo-advisor experience accumulates. The three clean placebos, anchoring (anc), extrapolation (ext), and confirmation (con), capture biases that do not require active position-formation and are predicted to be null. Additional bias measures (attention, disposition, gambling, herding) are reported in the Online Appendix. *Variable definitions.* τ is months since first robo-advisor adoption, capped at 24. Adopter equals one for investors who ever hold the robo-advisor. High Under-div. and High Turnover are indicators equal to one for adopters whose six-month pre-adoption mean of uhh (respectively tov) falls in the top decile of the eligible-adopter distribution; AOT (active opinion-taker) equals one if either high-decile indicator is one. Engagement is the investor's lagged 12-month login count, standardised. Never-adopters receive a pseudo-adoption date drawn from the empirical distribution of actual adopter adoption dates so that the calendar distribution of τ is comparable across treated and control samples. All specifications include one-month-lagged log of one plus total financial assets, lagged cash-management asset holdings, and lagged standardised login count as level controls; coefficients on these controls are omitted from the table for brevity.

Reading the falsification. The triple Adopter $\times\ \tau\ \times$ Engagement is the cleanest discriminator. It is negative and significant for the focal pair ($t = -3.07$ for uhh and $t = -6.68$ for tov) and statistically indistinguishable from zero for all three placebos. The within-adopter and DiD heterogeneity rows are numerically near-identical by construction, because the high-decile indicators are zero for never-adopters and the DiD triple interaction therefore reduces to the corresponding double interaction estimated on the adopter sample; both rows are reported because they arise from different estimating samples. Significance levels are based on standard-normal critical values applied to cluster-robust t-statistics, reported in parentheses below coefficient estimates: *** $|t| \geq 2.576$ ($p < 0.01$); ** $|t| \geq 1.960$ ($p < 0.05$); * $|t| \geq 1.645$ ($p < 0.10$); two-sided.

Panel B: Raw monthly bias (robustness)

	Focal: predicted decline		Placebo: predicted null		
Regressor	Under-div. (uhh)	Turnover (tov)	Anchoring (anc)	Extrap. (ext)	Confirm. (con)
Within-adopter sample					
τ	$+0.0001$	-0.0003	$+0.0007^{**}$	-0.0001	-0.0022^{***}
	$(+0.70)$	(-1.15)	$(+2.03)$	(-0.14)	(-2.75)
$\tau\times$ High Under-div.	-0.0058^{***}	-0.0011^{**}	$+0.0005$	$+0.0006$	-0.0007
	(-14.55)	(-2.19)	$(+0.97)$	$(+1.42)$	(-0.73)
$\tau\times$ High Turnover	-0.0011^{***}	-0.0062^{***}	$+0.0011^{*}$	-0.0002	$+0.0004$

Continued on next page

Table 6 (continued) – Panel B

	Focal: predicted decline		Placebo: predicted null		
Regressor	Under-div. (uhh)	Turnover (tov)	Anchoring (anc)	Extrap. (ext)	Confirm. (con)
	(−3.49)	(−11.94)	(+1.96)	(−0.40)	(+0.53)
$\tau\times$ AOT	-0.0028^{***}	-0.0038^{***}	+0.0006	+0.0001	+0.0003
	(−10.06)	(−8.94)	(+1.47)	(+0.24)	(+0.39)
Difference-in-differences with matched never-adopters					
Adopter $\times\ \tau$	+0.0002	-0.0006^{**}	$+0.0010^{***}$	$+0.0009^{*}$	−0.0020
	(+1.25)	(−2.55)	(+3.94)	(+1.85)	(−1.56)
Adopter $\times\ \tau\times$ Engagement	-0.0003^{***}	-0.0007^{***}	+0.0001	+0.0001	+0.0002
	(−4.90)	(−8.33)	(+0.86)	(+1.12)	(+1.19)
Adopter $\times\ \tau\times$ High Under-div.	-0.0059^{***}	-0.0012^{**}	+0.0004	+0.0005	−0.0009
	(−14.60)	(−2.39)	(+0.86)	(+1.23)	(−0.99)
Adopter $\times\ \tau\times$ High Turnover	-0.0011^{***}	-0.0062^{***}	$+0.0011^{**}$	−0.0002	+0.0004
	(−3.39)	(−11.95)	(+1.98)	(−0.39)	(+0.49)
Adopter $\times\ \tau\times$ AOT	-0.0028^{***}	-0.0039^{***}	+0.0006	+0.0001	+0.0002
	(−9.90)	(−8.97)	(+1.43)	(+0.17)	(+0.30)
Observations (adopter-only)	1,274,140	1,274,140	1,274,140	1,274,140	1,231,806
Observations (DiD sample)	1,781,292	1,781,292	1,781,292	1,781,292	1,722,690

Notes for Panel B. The dependent variable is the raw monthly bias measure as observed at month t, replacing the rolling six-month trailing mean used in Panel A. The estimator, sample, matched-control construction, variable definitions, controls, and inference are identical to Panel A; see the Panel A notes for full detail. *Reading the smoothing comparison.* The engagement dose-response (Adopter $\times\ \tau\times$ Engagement) and the focal-bias heterogeneity interactions retain their sign and statistical significance for both focal biases when the smoothing window is removed; for the focal pair the engagement coefficient becomes sharper in Panel B ($t = -4.90$ for uhh; $t = -8.33$ for tov) than in Panel A. The simple time term τ and the DiD baseline term Adopter $\times\ \tau$ for uhh, by contrast, lose significance in the raw monthly construction. The pattern is consistent with learning that operates heterogeneously across investors and is detectable in the smoothed series for the average adopter but only in subsamples for raw monthly snapshots. The clean placebos remain null on the engagement triple in both panels Significance levels follow Panel A: *** $|t| \geq 2.576$ ($p < 0.01$); ** $|t| \geq 1.960$ ($p < 0.05$); * $|t| \geq 1.645$ ($p < 0.10$); two-sided. Cluster-robust t-statistics are reported in parentheses below coefficient estimates.

Table 7: Return Predictors of Robo-Advisor Allocation Share

	Lag 1 month		Cumulative 3 months		Cumulative 6 months	
Variable	Probit ME	Outcome (RA_{share})	Probit ME	Outcome (RA_{share})	Probit ME	Outcome (RA_{share})
Panel A: Positive/negative decomposition						
A.1: Raw returns						
Fund return^{+}	-0.026^{***}	0.006^{***}	-0.029^{***}	0.007^{***}	-0.031^{***}	0.009^{***}
	(-14.711)	(2.708)	(-14.876)	(3.140)	(-16.606)	(3.228)
Fund return^{-}	-0.024^{***}	0.007^{***}	-0.024^{***}	0.006^{***}	-0.027^{***}	0.008^{***}
	(-9.180)	(4.954)	(-8.431)	(3.434)	(-8.098)	(3.406)
RA return^{+}		0.006^{*}		0.008^{***}		0.009^{***}
		(1.941)		(2.822)		(3.121)
RA return^{-}		0.001		0.005		0.006^{**}
		(0.361)		(1.625)		(2.373)
Inverse Mills ratio		0.119^{***}		0.119^{***}		0.098^{***}
		(2.631)		(3.347)		(4.241)
$\bar{N}$ (uncensored)		27,968		27,234		25,324
N (selection)	86,647		86,973		86,378	
Pseudo-R^2 / R^2	0.0589	0.1116	0.0596	0.1133	0.0596	0.1117
A.2: Excess returns over CSI 300						
Fund return^{+}	-0.022^{***}	0.008^{***}	-0.024^{***}	0.008^{***}	-0.026^{***}	0.011^{***}
	(-8.806)	(4.592)	(-9.418)	(6.203)	(-10.089)	(3.525)
Fund return^{-}	-0.024^{***}	0.012^{***}	-0.029^{***}	0.011^{***}	-0.029^{***}	0.010^{***}
	(-8.763)	(5.485)	(-9.280)	(5.521)	(-8.983)	(5.597)
RA return^{+}		-0.001		-0.004		-0.005
		(-0.256)		(-1.535)		(-1.499)
RA return^{-}		-0.002		-0.003		-0.002
		(-0.677)		(-1.467)		(-0.842)
Inverse Mills ratio		0.134^{***}		0.144^{***}		0.113^{***}
		(2.924)		(4.136)		(4.921)
$\bar{N}$ (uncensored)		28,247		26,914		25,312
N (selection)	86,691		86,507		86,166	
Pseudo-R^2 / R^2	0.0583	0.1115	0.0601	0.1138	0.0594	0.1115
Panel B: Above-median indicator						
B.1: Raw returns						
Fund: above median	0.002	-0.001	0.003	0.002	0.000	0.002
	(0.566)	(-0.176)	(0.542)	(0.532)	(0.044)	(0.411)
RA: above median		0.001		0.001		-0.003
		(0.195)		(0.375)		(-0.548)
Inverse Mills ratio		0.109^{**}		0.111^{***}		0.092^{***}
		(2.424)		(2.954)		(3.741)
$\bar{N}$ (uncensored)		28,008		27,234		25,324
N (selection)	86,719		86,973		86,378	
Pseudo-R^2 / R^2	0.0560	0.1016	0.0566	0.1033	0.0563	0.1022
B.2: Excess returns over CSI 300						
Fund: above median	0.002	-0.001	0.003	0.002	0.002	0.001
	(0.533)	(-0.182)	(0.670)	(0.495)	(0.220)	(0.321)
RA: above median		0.001		0.001		-0.001
		(0.220)		(0.173)		(-0.276)
Inverse Mills ratio		0.110^{**}		0.110^{***}		0.091^{***}
		(2.425)		(2.978)		(3.613)
N	86,515	27,995	86,762	27,221	86,166	25,312
Pseudo-R^2 / R^2	0.0560	0.1012	0.0566	0.1023	0.0563	0.1018
Baseline controls			Yes			
Province dummies			Yes			

Notes. Stage-1 coefficients are reported as approximate average marginal effects on the robo-advisor adoption probability, $\mathrm{ME}_k = \phi(\Phi^{-1}(\bar{P}))\hat{\beta}_k$ with $\bar{P} \approx 0.323$ (average monthly RA-active rate) and scaling factor 0.359; t-statistics and significance are invariant to the scaling. Estimation uses Fama and MacBeth (1973) monthly cross-sectional Heckman two-step regressions; the reported coefficient is the time-series mean of the monthly estimates and the t-statistic is the ratio of that mean to its Newey–West HAC standard error with lag truncation $\min(3, T-1)$. The within R^2 is the plain R^2 averaged across monthly Stage 2 regressions; the pseudo-R^2 is the McFadden pseudo-R^2 averaged across monthly Stage 1 probits. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed).

Table 8: Pre-RA Experience and Robo-Advisor Allocation over Tenure

Variable	Interaction: linear OLS	Interaction: quadratic OLS	Interaction: quadratic Fractional logit
Pre-RA experience and tenure interactions			
Pre-RA experience	-0.196***	-0.196***	-1.318***
	(-2.622)	(-2.884)	(-3.530)
Pre-RA experience × tenure	-0.023*	-0.043**	-0.304**
	(-1.912)	(-1.972)	(-2.247)
Pre-RA experience × tenure^2	—	-0.013	-0.071
		(-1.232)	(-1.248)
Tenure (main effects)			
Tenure (within-month demeaned)	-0.028**	-0.048**	-0.330**
	(-2.194)	(-2.109)	(-2.394)
Tenure^2	-0.000	-0.013	-0.073
	(-0.390)	(-1.267)	(-1.268)
Province Dummies	Yes	Yes	Yes
N	1,097,096	1,097,096	1,097,096
R^2	0.162	0.163	0.190

Notes. The sample consists of investor-months in which the investor holds positive robo-advisor assets (the adopter-only subsample). The dependent variable $rsh_{i,t} = ra_assets_{i,t}/(ra_assets_{i,t} + mf_assets_{i,t})$ is the share of risky-asset wallet allocated to the robo-advisor and is bounded in $[0, 1]$. Because the analysis is conducted within the adopter sample, the Heckman selection correction used elsewhere in the paper is not applied here.

Pre-RA experience is the investor-level average of *exp_fund_months* over the $L_{\text{PRE}} = 6$ months immediately before the investor's first robo-advisor month, requiring at least three non-missing months; the quantity is z-scored once across eligible adopters and is therefore time-invariant for each investor. *Tenure* is the cumulative count of robo-advisor- active months for the investor (gap months do not accrue), demeaned within each calendar month to absorb time effects; *Tenure*^{2} is the demeaned tenure squared. The interaction *Pre-RA experience* × *Tenure* is the focal regressor; the reclaiming-control hypothesis predicts a negative coefficient on this term, meaning that adopters with more pre-RA fund-investing experience reduce their robo-advisor share of wallet as robo-advisor tenure accumulates.

Estimation uses Fama and MacBeth (1973) monthly cross-sectional regressions. The reported coefficient is the time-series mean of the monthly cross-sectional estimates; the reported t-statistic is the ratio of that mean to its Newey-West HAC standard error with lag truncation $\min(3, T-1)$. All three columns include *Tenure* and *Tenure*^{2} as main effects. The columns differ in the skill-tenure interaction specification. Column (1) includes only the linear interaction *Pre-RA experience* × *Tenure*; columns (2) and (3) include both the linear interaction and the quadratic interaction *Pre-RA experience* × *Tenure*^{2}. Columns (1) and (2) are estimated by ordinary least squares; column (3) replaces OLS with a fractional-logit generalised linear model (binomial family, logit link; Papke and Wooldridge, 1996) to accommodate the bounded support of the dependent variable. Coefficients in column (3) are on the logit scale and are not directly comparable in magnitude to the OLS coefficients in columns (1) and (2). The R^2 reported for columns (1) and (2) is the plain R^2 averaged over monthly cross-sectional regressions; for column (3) it is the McFadden pseudo-R^2 ($1 - \text{deviance}/\text{null deviance}$) averaged across months and is not comparable to the OLS values. The total-observations row reports the sum of monthly cross-sectional sample sizes across the $T = 44$ months in which a specification clears the pre-flight checks (degrees of freedom, finite design matrix, full column rank, condition number $< 10^{12}$, and GLM convergence for column (3)).

Baseline controls (log total financial assets$_{t-1}$, stated risk preference, age, female indicator, combined-advisor dummy$_{t-1}$, cash-management assets$_{t-1}$, months-since-registration$_{t-1}$, and province dummies) are included in every regression but suppressed for brevity; full estimates are reported in the Online Appendix.

Newey-West HAC t-statistics in parentheses below coefficient estimates. Significance levels (against standard-normal critical values for the Fama-MacBeth time-series mean): *** $|t| \geq 2.576$ ($p < 0.01$); ** $|t| \geq 1.960$ ($p < 0.05$); * $|t| \geq 1.645$ ($p < 0.10$); two-sided.

Table 9: Robo-Advisor Dose-Response on the Self-Directed Mutual-Fund Sleeve

	(1) Holds MF (LPM)	(2) Concentration (HHI)	(3) No. funds	(4) Stock share	(5) Risky share	(6) Turnover
RA-share bins at $t-1$ (Bin 0: $rsh_{t-1}=0$ omitted as reference)						
Bin 1 (0–0.20]	0.005***	-0.216***	0.196***	-0.001	-0.001	-0.105***
	(7.188)	(-50.801)	(27.425)	(-0.342)	(-0.148)	(-18.416)
Bin 2 (0.20–0.40]	0.005***	-0.182***	0.052***	0.025***	0.018*	-0.035***
	(7.496)	(-30.097)	(13.509)	(3.551)	(1.943)	(-6.684)
Bin 3 (0.40–0.60]	0.000	-0.023***	-0.026***	0.032***	0.013	0.117***
	(0.360)	(-2.943)	(-5.953)	(2.653)	(0.854)	(14.131)
Bin 4 (0.60–0.80]	-0.016***	0.117***	-0.084***	0.060***	0.016	0.299***
	(-9.199)	(11.166)	(-13.206)	(4.173)	(1.001)	(23.017)
Bin 5 (0.80–1.00]	-0.332***	0.396***	-0.122***	0.008	-0.170***	0.805***
	(-19.380)	(26.236)	(-16.846)	(0.462)	(-6.842)	(36.967)
log(TFA), $t-1$	0.015***	-0.454***	0.261***	-0.071***	0.108***	-0.549***
	(18.403)	(-34.248)	(23.040)	(-10.161)	(10.708)	(-61.998)
Baby assets, $t-1$	-0.002***	0.061***	-0.035***	-0.012***	-0.039***	0.124***
	(-11.430)	(30.271)	(-22.460)	(-4.065)	(-17.600)	(49.691)
Combined-advisor dummy, $t-1$	-0.001	-0.080***	0.187***	-0.026***	-0.009	-0.035***
	(-1.281)	(-12.568)	(19.274)	(-3.596)	(-1.258)	(-6.584)
Investor FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Sample	Full	MF active	MF active	MF active	MF active	MF active
N	4,492,147	4,325,040	4,325,040	4,339,137	4,339,137	4,075,169
Investors	100,000	100,000	100,000	100,000	100,000	98,067
Within R^2	0.102	0.125	0.114	0.003	0.004	0.108

Notes. The table reports two-way fixed-effects panel regressions (`PanelOLS` with investor and month effects) of self-directed mutual-fund (MF) sleeve outcomes on lag-1 robo-advisor (RA) intensity. RA intensity is $rsh_{i,t-1} = RA_{i,t-1}/(RA_{i,t-1} + MF_{i,t-1})$, binned into six categories; the no-RA baseline (Bin 0, $rsh_{t-1}=0$, which includes never-adopters and pre-adoption months, 65.0% of rows) is the omitted reference. Each bin coefficient is the within-investor change in the period-t outcome relative to Bin 0. The five non-zero bins hold 26.8%, 4.1%, 1.8%, 0.9%, and 1.4% of panel rows respectively. Outcomes: column (1) is a linear probability model for holding any mutual fund (*mf_active*); columns (2)–(6) are estimated on the MF-active subsample. Concentration is the within-MF Herfindahl index; No. funds is the count of distinct funds held; Stock share is (stock + index + mixed) over total MF; Risky share is one minus the money-market fraction; Turnover is the active-trading proxy. Outcomes (2)–(6) are winsorised at 1/99 and z-scored within month, so coefficients are in within-month standard-deviation units. Controls (log total financial assets$_{t-1}$, cash-management assets$_{t-1}$, both winsorised and z-scored within month, and a combined-advisor dummy$_{t-1}$) are reported. Investor demographics, skills, and province are absorbed by investor fixed effects. Standard errors are two-way clustered on investor and month; clustered t-statistics in parentheses. *Interpretation.* The treatment is lagged one month, which guards against contemporaneous reverse causation but not joint within-month determination of RA intensity and MF behaviour. The estimates describe within-investor co-movement across the RA-share dose, not a causal effect of RA usage. Secondary outcomes (auto-invest share, gambling bias, disposition effect) are reported in the Online Appendix. Stars: *** $p<0.01$, ** $p<0.05$, * $p<0.10$ (two-tailed).

Table 10: Return Predictors of Robo-Advisor Allocation across Adoption-Lifecycle Margins

	Lag 1 month		Cumulative 3 months		Cumulative 6 months	
Variable	Probit ME	Outcome (RA_{share})	Probit ME	Outcome (RA_{share})	Probit ME	Outcome (RA_{share})
Panel A: First adoption						
Fund return$^+$	-0.003***	0.004	-0.003***	0.009	-0.002***	0.024**
	(-7.014)	(0.525)	(-5.572)	(1.574)	(-5.725)	(2.058)
Fund return$^-$	-0.006***	-0.001	-0.006***	0.014	-0.006***	0.042**
	(-5.878)	(-0.046)	(-5.765)	(1.254)	(-5.864)	(2.354)
RA return$^+$						
RA return$^-$						
Inverse Mills ratio		0.096		-0.124		-0.432**
		(0.464)		(-1.029)		(-2.433)
N	1,704,054	54,568	1,717,153	48,672	1,715,065	46,920
Pseudo-R^2 / R^2	0.033	0.184	0.030	0.188	0.030	0.193
Panel B: Exit (probit-only design; marginal effects)						
Fund return$^+$	0.001***		0.002***		0.003***	
	(2.783)		(4.703)		(8.644)	
Fund return$^-$	-0.002***		-0.002***		-0.001**	
	(-3.017)		(-2.912)		(-2.306)	
RA return$^+$	0.000		-0.000		0.000	
	(0.777)		(-0.562)		(0.262)	
RA return$^-$	0.002***		0.001		0.001	
	(2.928)		(0.688)		(0.986)	
N (selection)	1,263,735		1,177,032		1,118,619	
Pseudo-R^2	0.025		0.026		0.026	
Panel C: Continuation						
Fund return$^+$	-0.003**	0.012***	-0.004***	0.012***	-0.006***	0.014***
	(-2.375)	(10.882)	(-3.291)	(8.562)	(-8.449)	(8.265)
Fund return$^-$	0.004***	0.012***	0.004***	0.012***	0.003**	0.012***
	(3.130)	(10.702)	(3.024)	(9.150)	(2.290)	(6.168)
RA return$^+$		0.007**		0.009***		0.009***
		(2.322)		(2.931)		(3.634)
RA return$^-$		0.002		0.005		0.007**
		(0.501)		(1.629)		(2.370)
Inverse Mills ratio		0.445***		0.539***		0.411**
		(3.296)		(3.984)		(2.511)
N	1,412,944	1,286,824	1,444,456	1,316,070	1,406,508	1,124,192
Pseudo-R^2 / R^2	0.020	0.108	0.021	0.111	0.020	0.110
Baseline controls	Yes	Yes	Yes	Yes	Yes	Yes
Province dummies	Yes	Yes	Yes	Yes	Yes	Yes

Notes. Estimates are from unweighted Heckman two-step (Tobit-2) models estimated monthly and aggregated by Fama–MacBeth. The Stage 1 selection equation is a probit on the relevant participation margin; the Stage 2 outcome equation is an OLS of the robo-advisor allocation share $ra_share_{i,t} = RA_{i,t}/(RA_{i,t} + MF_{i,t})$ on the uncensored subsample with the inverse Mills ratio included. The selection margin differs by panel: in Panel A, $D_{i,t} = 1$ in the investor's first robo-advisor month and 0 in all prior months (post-adoption months excluded); in Panel B (probit-only), $D_{i,t} = 1$ if the investor held RA in $t-1$ and not in t; in Panel C, $D_{i,t} = 1$ if the investor held RA in both $t-1$ and t, 0 if the investor held in $t-1$ but exited in t. Each return r is decomposed into its positive part $r^+ = \max(r, 0)$ and negative part $r^- = \max(-r, 0)$, both winsorised at 1/99 and z-scored within month, and entered jointly. The fund (Fund) and robo-advisor (RA) raw monthly returns enter the same regression. All Stage-1 coefficients are reported as approximate average marginal effects on the relevant participation probability, $\text{ME}_k = \phi(\Phi^{-1}(\bar{P}))\hat{\beta}_k$, where $\bar{P}$ is the average monthly selection rate in each design and $\hat{\beta}_k$ is the Fama–MacBeth average probit coefficient; t-statistics and significance are invariant to this scaling. The scaling factor is 0.072 in Panel A (first adoption, $\bar{P} \approx 0.032$), 0.061 in Panel B (exit, $\bar{P} \approx 0.027$), and 0.161 in Panel C (continuation, $\bar{P} \approx 0.911$). Because continuation is a near-ceiling event ($\bar{P}$ close to one), its Panel C marginal effects describe the response of the marginal would-be exiter and should be read with that caveat; the underlying probit coefficients are correspondingly larger. Baseline controls (log total financial assets$_{t-1}$, stated risk preference, age, female indicator, combined-advisor dummy$_{t-1}$, cash-management assets$_{t-1}$, and province dummies) are included and suppressed in both stages. Months-since-registration ($t-1$) enters Stage 1 only and is the exclusion restriction. $\bar{N}$ (uncensored) is the average number of RA-adopter months per cross-section in Stage 2; N (selection) is the total number of at-risk investor-months in Stage 1. In the first-adoption (Panel A) and continuation (Panel C) designs the recovered $\hat{\rho} = \hat{\lambda}/\hat{\sigma}$ exceeds one in magnitude in several months, because $\hat{\lambda}$ is large relative to the within-month residual standard deviation; the inverse Mills ratio should be read as a Heckman selection diagnostic rather than a strict correlation in those panels. Newey–West HAC t-statistics (lag = $\min(3, T-1)$) in parentheses. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed).

A. Online Appendix

This appendix collects supplementary material referenced in the main text. It serves three purposes. First, it documents sample construction, estimation choices, and variable construction details that are too long for the main text but are necessary for evaluating the empirical design. Second, it reports robustness checks for the two-margin Heckman estimates, the authority ladder, the reclaiming-control specification, performance predictors, and the post-adoption bias-trajectory evidence. Third, it records auxiliary evidence on external validity, age-bias correlations, and the structure of the self-directed sleeve. The appendix is descriptive and interpretive in the same sense as the main text: the estimates document revealed demand patterns and mechanism-consistent heterogeneity, but they are not presented as causal treatment effects of robo-advising.

A.1. External Validity and Single-Platform Representativeness

The analysis is conducted on a single platform in one country, which raises a natural external-validity concern. We do not use the Alipay setting to estimate a universal adoption level for robo-advising. The levels of adoption, the age profile of users, and the role of platform discovery are likely to differ in Western brokerage markets, retirement-plan markets, and advisory markets. The portable object is instead the two-margin logic. Delegated financial products can have fixed entry costs and marginal authority costs, the same investor characteristic can move the adoption margin and the conditional-delegation margin in opposite directions, and persistence responds to its own forces. Three institutional features make the setting useful for that purpose. First, Alipay is a large mobile payments and wealth-management platform in China, and its account structure records mutual-fund holdings, delegated robo-advisor assets, automatic investment plan activity, cash-management balances, platform engagement, transactions, and returns in one place. This integrated observation reduces the external-holding problem that would arise in a single-brokerage dataset where the researcher observes only one slice of the investor's portfolio. Second, BangNiTou is a fully delegated mutual fund robo-advisor that constructs and rebalances portfolios for the investor. Its economic content is similar to delegated portfolio products in other markets even though its institutional setting is Chinese. Third, automatic investment plans coexist in the same account. They automate scheduled fund purchases but leave fund choice with the investor. The coexistence of self-directed funds, automatic investment plans, and robo-advising is what allows the paper to distinguish execution automation from delegation of fund-selection authority. The limitations are equally important. The Alipay investor base is younger than many Western brokerage populations; the Chinese mutual fund market has its own regulatory and distribution structure; and the super-app environment makes platform tenure and product discovery more central than they may be in a fragmented brokerage environment. We therefore interpret demographic magnitudes and adoption levels as context-specific. The stronger claim is about signs, cross-margin differences, and the authority-versus-automation distinction, not about the level of robo-advisor participation in other markets.

A.2. Additional Methodological and Estimation Details

The main Heckman estimates are run month by month and then averaged across months. In each monthly cross-section, the first stage is a probit for robo-advisor participation and the second stage is an ordinary least squares regression of the conditional robo-advised share on observed covariates and the inverse Mills ratio. The reported coefficients are Fama–MacBeth averages, and inference uses Newey–West HAC standard errors with maximum lag three. This design keeps the identifying variation cross-sectional, which is appropriate for investor characteristics that move slowly over time, while allowing the slope coefficients to vary across market environments. The monthly estimation pipeline applies several safeguards. We drop regressors that are constant within month, verify the rank of the probit and outcome design matrices, remove variables that perfectly separate the adoption outcome, drop the joint set of province dummies when the effective parameter count approaches the number of adopters, and skip monthly cross-sections in which the probit fails to converge. Province dummies are included as controls when the monthly design matrix supports them, but their individual coefficients are not reported. These guards are mechanical and are applied before interpreting any coefficient. The number of retained months differs across specifications because richer specifications require more within-month variation and because behavioral measures rely on rolling windows that are not populated in the earliest months. The baseline and skills specifications retain 48 and 47 monthly cross-sections, and the bias-augmented specification 36. All continuous regressors are winsorized within month at the first and ninety-ninth percentiles and then standardized within month. Binary indicators are not standardized. Except for variables that are time-invariant or near-invariant, covariates are lagged one month so that they are predetermined relative to the month-t participation and allocation decisions. Returns are never forward filled. Behavioral measures are forward filled across non-trading months only after the investor's first observed trading month, reflecting the premise that a measured behavioral tendency remains informative until new trading behavior updates it. Selection-stage coefficients in tables without interactions are rescaled to approximate marginal effects using the observed event rate for the relevant sample,

$$\mathrm{ME}_k = \phi\big(\Phi^{-1}(\bar{P})\big)\,\hat{\gamma}_k, \tag{15}$$

where $\bar{P}$ is the relevant monthly event rate. The main Heckman uses panel-specific RA-active rates of 0.326 in the baseline panel, 0.344 in the skills panel, and 0.372 in the biases panel, giving scaling factors of 0.360, 0.368, and 0.378. Outcome-stage coefficients are reported in units of the delegated share. In tables with interactions, the selection-stage coefficients are left on their native scale because the marginal effect of an interaction in a probit is not the coefficient on the interaction term and depends on the full covariate vector (Ai & Norton, 2003). The authority-ladder specifications avoid this issue by using a stacked linear probability design with investor-month fixed effects.

A.3. Choice-Based Sampling and Unweighted Estimation

The main sample is stratified on ever adoption. The purpose of the stratification is statistical power, not prevalence estimation. First adoption and re-entry are rare monthly events, and the conditional delegation share is observed only when the investor holds robo-advised assets. A uniform random

sample of the same size would deliver many non-user months and too few robo-active observations for precise estimation of the conditional-share equation, lifecycle margins, and post-adoption dynamics. The main estimates are unweighted because the object of interest is the slope relationship between investor characteristics and the two margins, conditional on the active mutual-fund investor population, not the platform-level prevalence of robo-advisor use. Population-calibration weights would rescale ever-adopters and never-adopters back toward their population frequencies, but in thin monthly cross-sections they also amplify the influence of rare observations and increase the variance of the Fama–MacBeth slopes. We therefore read the unweighted coefficients as describing the sampled population, and report the weighted estimates as the check on whether the sign pattern is specific to it. We do not assert that the unweighted slopes are consistent for the platform population; establishing that would require a retrospective-sampling argument for a probit selection stage with full-window stratification, which we do not derive. We nevertheless re-estimate the main Heckman specifications with calibrated inverse-probability weights in Table A.1. Within each month, adopter observations receive weight $p_{pop}/p_{samp,t}$ and non-adopter observations $(1-p_{pop})/(1-p_{samp,t})$, normalized to mean one, with $p_{pop} = 1/101$. The headline sign patterns are unchanged: fund experience remains a positive participation predictor with a negative conditional-share relation; underdiversification and turnover preserve the negative-adoption, positive-delegation signature; and confirmation preserves the opposite pair of signs. Because the weights center the fitted adoption probability on the population rate, the reported selection marginal effects are evaluated at a one-percent base rate and are not comparable in magnitude with the main table; significance should be read from the t-statistics. The inverse Mills coefficient turns negative under weighting, and the selection-on-unobservables discussion in the main text therefore rests on the unweighted panel. The weighted estimates support the conclusion that the main sign findings are not artifacts of the unweighted stratified design.

A.4. Variable Construction Details

This section records the construction of the behavioral and skill measures used in the main text. The measures are designed to capture distinct investment patterns, not a single behavioral type. Whenever possible, the measure is constructed from information available before the decision month and avoids mechanical overlap with the outcome being explained. Underdiversification is measured primarily by the Herfindahl index of mutual fund portfolio weights,

$$\text{HHI}_{it} = \sum_{f} w_{ift}^2,$$

where w_{ift} is investor i's end-of-month weight in fund f. Higher HHI means a more concentrated fund portfolio. We also compute related diversification measures for robustness, including the effective number of funds, $1/\text{HHI}_{it}$, pairwise return-correlation measures, and a risk-adjusted diversification ratio. HHI is the main text measure because it is defined whenever the investor holds at least one mutual fund, whereas correlation-based measures require multiple holdings with sufficient return history. Total turnover is an active-trading intensity measure. It is computed from monthly buys and sells relative to average holdings between $t-1$ and t, using a Modified-Dietz-style denominator. Speculative turnover isolates a narrower form of active trading. A trade is classified as speculative

when the investor held the fund at $t-1$, the position had a positive lagged one-month return, the investor reduced the position in month t, and the investor repurchased either the same fund or a fund of the same type within one month. The same-type version is the primary measure because same-fund round trips are rare in mutual fund data; the strict same-fund version is used as a robustness variant. Confirmation bias is the capped retention rate of lagged winner positions. We first identify positions held at the end of $t-1$ whose trailing twelve-month return was positive. For each such position, retention is the fraction of the $t-1$ value still held at t, capped at one so that new purchases do not inflate the measure. The investor-level confirmation-bias measure is the value-weighted average retention rate across lagged winners. Disconfirmation bias is constructed analogously from lagged loser positions as an exit rate. Because winners and losers are disjoint sets, the two measures capture different realized-holding responses. The disposition effect follows the standard PGR-minus-PLR logic. PGR is the proportion of gains realized and PLR is the proportion of losses realized. The position panel is rooted in holdings observed at $t-1$ and then follows the investor's decision during month t. This timing matters because complete exits between $t-1$ and t are invisible if the measure is constructed only from end-of-month-t holdings. Attention bias is the share of trading value directed to salient funds. Salience flags include funds with extreme trailing twelve-month returns, funds with top-decile current-month returns within style or across the fund universe, and newly launched funds. We exclude top-Sharpe and low-volatility flags from the primary measure because those variables can reflect informed attention rather than limited attention. Extrapolation bias is the share of trading value directed to funds whose lagged one-month return lies in the top decile, computed within Morningstar style and across the fund universe and combined through the union of the two criteria. The lagged construction is important because contemporaneous returns can be mechanically affected by flows. Gambling bias adapts the lottery-stock logic to mutual funds. A fund is lottery-like when it has low net asset value, high idiosyncratic volatility, and high idiosyncratic skewness, measured using lagged return information. Because diversified mutual funds rarely satisfy all three criteria at once, the primary measure uses an any-two-of-three rule and the strict triple-condition version is retained for robustness. Anchoring captures trading near prior price extremes. Anchored sells are sales of funds whose lagged price is close to the previous twelve-month high. Anchored buys are purchases of funds whose lagged price is close to the previous twelve-month low. The two directional ratios are combined as their mean in the main measure. Herding is the share of the investor's trading value directed to funds with high lagged aggregate flows. Aggregate flow is computed with a leave-one-out correction so that investor i's own trading does not mechanically enter the fund-level flow signal used to classify her behavior. Selection ability and timing ability decompose realized mutual-fund performance into two pieces. Selection ability is the value-weighted excess return of the investor's held funds relative to Morningstar category benchmarks. Timing ability captures active profit and loss associated with within-month position changes, orthogonal to the passive return on beginning-of-month holdings. Fund experience is the cumulative number of months in which the investor has active mutual fund holdings. Experience is therefore a stock of market participation, not a realized-performance measure.

A.5. Full Specification and Alternative Intensive Outcomes

Online Appendix Table A.3 reports the joint Heckman specification with demographics, skills, and all behavioral variables included simultaneously. This specification is not the main table because the rolling-window construction of the behavioral variables restricts the retained month sample and because the nested specifications in the main text make clear which block absorbs which relation. The joint table is nevertheless the most direct check that the headline sign patterns are not artifacts of estimating skills and biases in separate panels. The three strongest behavioral sign patterns survive jointly: underdiversification remains negative on adoption, at -0.065, and positive on conditional delegation, at $+0.016$; turnover remains negative on adoption, at -0.028, and positive on conditional delegation, at $+0.039$; and confirmation remains positive on adoption, at $+0.019$, and negative on conditional delegation, at -0.011. Fund experience retains a positive participation coefficient, at $+0.024$, but its conditional-share coefficient attenuates to -0.005 once biases are included jointly. The interpretation in the main text follows from this pattern. Experience is robustly related to participation, while part of its conditional-share relation is mediated by behavioral correlates of experience. The inverse Mills ratio remains positive in the joint specification, at $+0.277$, so the paper treats selection on unobservables as an economic object rather than a nuisance term fully absorbed by the observed controls. Online Appendix Table A.4 reports alternative intensive-margin outcomes. Log RA assets capture absolute delegated scale, while RA assets over total financial assets capture delegation relative to the full on-platform balance sheet. These alternatives check whether the main results are mechanical consequences of the RA-plus-mutual-fund denominator in the main outcome. The core sign patterns are preserved across the alternative outcomes: turnover and underdiversification are positive on both alternatives, and confirmation negative on both. The inverse Mills coefficient changes across outcomes, however. It is positive on the main compositional share, smaller and only marginally significant on the total-asset share, and strongly negative on the log-level outcome. The pattern locates the channel, since the positive selection channel is strongest when delegation is measured as a composition of the investor's fund holdings instead of as the absolute level of assets inside the robo-advisor.

A.6. Panel Fixed-Effects Control-Function Estimates

The main estimates use a month-by-month Heckman with Fama–MacBeth aggregation and Newey–West inference. A natural question is whether the behavioral cross-margin signs depend on that estimator or on its cross-sectional identification. Online Appendix Table A.5 re-estimates the conditional-share equation with a completely different design: the control-function panel selection estimator of Wooldridge (1995) and Semykina and Wooldridge (2010). The first stage forms the inverse Mills ratio from per-month probits augmented with Mundlak per-investor time averages and the months-since-registration exclusion restriction, and the second stage is a linear panel regression that adds the inverse Mills ratio to investor fixed effects, in one variant, and to investor and month fixed effects, in the other, with two-way clustered standard errors. This swaps Fama–MacBeth cross-sectional averaging for within-investor identification and swaps Newey–West for two-way clustering, so agreement across the two estimators bears directly on whether the signs are an artifact of either choice. The three focal intensive-margin behavioral signs are preserved. In the biases specification, total turnover and

underdiversification carry positive and highly significant coefficients on the conditional share, 0.040 and 0.045 with investor fixed effects and 0.041 and 0.032 with investor and month fixed effects, and confirmation bias is negative and highly significant, near -0.015 in both variants. Panel B repeats the exercise on the risky-RA-share outcome and recovers the same three signs at similar magnitudes, 0.045 and 0.046 for turnover, 0.052 and 0.038 for underdiversification, and -0.015 for confirmation, so the cross-margin behavioral pattern is not specific to the main denominator. The fixed-effects design cannot speak to the demographic and risk-preference gradients, because risk preference, age, and gender are time-invariant within investor and are absorbed by the entity fixed effect, so the table is read as a robustness check on the behavioral intensive-margin signs rather than on the extensive margin or the demographic cross-section. The inverse Mills coefficient is not stably signed across specifications in this design, and because the second stage treats the inverse Mills ratio as observed rather than as a generated regressor, its standard errors understate uncertainty and its t-statistics should be read as approximate. We therefore do not use the panel control-function selection term to corroborate the positive selection on unobservables documented in the main Fama–MacBeth tables; the contribution of the table is narrower and cleaner, namely that the behavioral cross-margin signs on the conditional share survive a fixed-effects estimator and two-way clustered inference, and are not produced by the estimator and the inference method of the main specification.

A.7. AUTO, Authority, and Reclaiming-Control Robustness

Online Appendix Table A.6 reports a stand-alone AUTO Heckman, parallel to the main RA Heckman. AUTO denotes automatic investment plans, which automate purchase timing while preserving investor control over fund selection. The stand-alone AUTO estimates show that AUTO and RA are not the same product at different intensities. Underdiversification is strongly positive on AUTO intensity, at $+0.088$ against $+0.020$ for RA; confirmation is positive on AUTO entry and positive, not negative, on AUTO intensity; turnover is negative on both AUTO margins where it is positive on the RA conditional share; and the AUTO inverse Mills coefficient is negative in the biases specification. The opposite sign of selection on unobservables relative to RA is especially useful for interpretation, since it indicates that latent demand for intensive AUTO use is not the same residual object as latent demand for intensive RA use. Online Appendix Table A.9 reports the authority-ladder robustness checks and the baseline-control interactions suppressed in the main table. The first three rows of each block vary the pre-use horizon over which investor type is measured to three, nine, and twelve months, holding the identifying-only restriction fixed; the fourth row holds the horizon at the headline six months and instead retains all eligible investor-months. The distinction between the automation step and the full-delegation step is preserved throughout. Engagement, underdiversification, anchoring, and attention reject equality of the two steps at every horizon and in the full panel; total turnover rejects at every horizon the table reports, its $L = 3$ block having not been supplied. The skill composite, marginal on the full-delegation step in the headline six-month window, reaches $+0.016$ at nine and twelve months while its automation-step coefficient turns negative, though its equality test never rejects at five percent. Selection ability is already significant at five percent in the headline window, holds at $+0.013$ at nine months, and weakens to $+0.012$ and ten percent significance at twelve. The sample shrinks as the horizon lengthens, from 766,254 stacked observations at $L = 3$ to 496,575 at

$L = 12$, so comparisons across horizons are across partly different investors, and we read the horizon ladder as a stability check rather than as independent evidence. Online Appendix Table A.10 reports reclaiming-control robustness under three alternative adopter-sample restrictions: consecutive RA-months only, no-gap stayers, and dropping the first RA-active month per investor. The baseline specification studies RA-active adopter-months and relates delegated share to RA tenure, pre-RA fund experience, and their interaction. Across these variants, the pre-RA-experience by tenure interaction remains negative, at -0.036, -0.040, and -0.036, and the tenure main effect remains negative at -0.042 to -0.045. The result should be read as a tenure-share gradient among active adopter months, not as a randomized effect of experience, but its stability supports the interpretation that experienced investors use delegation less persistently.

A.8. Selective Awareness and Bias-Trajectory Robustness

Online Appendix Table A.11 reports additional bias-trajectory specifications over robo-advisor tenure for attention, the disposition effect, gambling, and herding. These specifications complement the main selective-awareness evidence. The main text emphasizes the behaviors most directly tied to corrective portfolio management, underdiversification and turnover. The appendix expands the trajectory evidence to related behavioral measures and interaction variants. The purpose is not to claim that adoption causally eliminates biases. The purpose is to ask whether the behaviors that generate the negative-adoption, positive-delegation signature move after exposure in a direction consistent with learning from, or adjusting to, the delegated portfolio. The evidence is strongest for underdiversification and turnover. These behaviors decline with robo-advisor tenure, and the decline is more pronounced among investors with high pre-adoption bias or high platform engagement. Engagement is interpreted cautiously. It may raise exposure to portfolio composition and therefore salience, but it may also proxy for sophistication or general financial attention. The appendix therefore treats the interaction as mechanism-consistent heterogeneity, not as unique identification of awareness. The four additional biases behave in ways that support, but do not add to, the focal reading. Attention is effectively null: its adopter-by-tenure term is indistinguishable from zero in both panels, and its engagement triple is significant in the smoothed panel but shrinks to insignificance without the smoothing window, which is what a smoothing artifact looks like. Gambling is not null in the matched design, where the adopter-by-tenure term is -0.0006 and -0.0008 and significant in both panels, though it is flat within adopters; its engagement triple is positive, so it does not show the focal pattern in any case. The disposition effect declines with tenure in both panels, but the interpretation is mechanical not psychological. The measure is built from investor-initiated realizations on the self-directed book, and delegation moves wealth into a sleeve the investor does not trade, so the denominator of the statistic falls with tenure. Herding rises with tenure in the smoothed panel, at $+0.0004$ in the matched design, which is the opposite of the focal prediction, though the raw monthly counterpart is insignificant and slightly negative; the natural reading of the smoothed result is a tracking channel, in which the investor observes the robo-advisor's portfolio adjustments and partially copies them in her own book. None of the four is a clean placebo, which is why the main table uses anchoring, extrapolation, and confirmation instead. The appendix also clarifies the relation between selective awareness and the other behavioral channels. Selective awareness is the paper-level mechanism for correctable

active-opinion mistakes. Attachment and cognitive constraints are signatures inside the behavioral taxonomy. Confirmation-bias patterns are read through attachment to validated positions or narratives. Attention, gambling, and herding are read as constraints or preferences that robo-advising does not directly relax. The trajectory evidence is not expected to be uniform across all biases, because the framework predicts different demand signatures for different behavioral mechanisms.

A.9. Performance Appendices

Online Appendix Table A.7 adds single fund-performance measures to the RA Heckman. The outcome-stage coefficients are economically very small, with absolute values no larger than 0.003. Category-adjusted fund returns at longer horizons are statistically significant but too small to explain the main cross-margin patterns. The result matters because a pure performance-chasing account would predict recent self-directed fund performance to explain a large part of delegation intensity. It does not. The inverse Mills ratio in these specifications remains positive and significant for the excess-return and category-adjusted arms, between 0.087 and 0.098, consistent with the main tables. Online Appendix Table A.8 adds RA-performance measures. Broad RA excess returns and beta-adjusted returns are small, no larger than 0.005 in absolute value at any horizon, while strategy-adjusted RA returns are positive, at 0.007 to 0.008. The main text interprets this as within-RA strategy selection rather than broad evidence that investors increase RA allocation because the RA channel beats the market. This distinction is important because investors may choose among risk profiles or strategy sleeves within RA even when broad RA performance does not explain whether they delegate more in the first place. Online Appendix Tables A.12, A.13, and A.14 extend the performance analysis to first adoption and lifecycle margins. Before first RA adoption, investor-specific RA returns are mechanically undefined, so the relevant observed signal is recent self-directed fund performance and broader platform-level return information. After adoption, both the self-directed and delegated sleeves can generate realized return signals. The lifecycle return tables therefore separate first adoption, continuation, exit-related prediction, and conditional allocation. Across these exercises, performance feedback matters in some decomposed specifications, especially when positive and negative components are separated, but it does not absorb the main characteristic-based cross-margin patterns. The first-adoption return tables are estimated on short monthly panels, fifteen to twenty-seven cross-sections, and their coefficients take opposite signs between the pre-window and lag-1 hazard constructions, so we treat them as evidence of absence of a strong return channel, not as precise estimates.

A.10. Dose-Response and Self-Directed-Sleeve Outcomes

Online Appendix Table A.15 complements the main dose-response table by reporting three secondary outcomes: the auto-invest share of fund purchases, gambling bias, and the disposition effect. The omitted group is no lagged RA share, so the coefficients compare investor-months with different lagged RA-share bins to investor-months without active RA exposure, conditional on the fixed effects and controls specified in the table. The auto-invest share rises at low RA intensity, by 0.064 within-month standard deviations in the lowest bin, and turns sharply negative at the top bin, at -0.103; the outcome is z-scored within month, so these are not percentage points. Gambling bias and the disposition effect, by contrast, are essentially unmoved by RA intensity: their within-R^2 is zero to three

decimals and almost every bin coefficient is insignificant. They function as placebo tests. The restructuring of the sleeve documented in the main table, across substitution, concentration, turnover, and risk, is not accompanied by a shift in these underlying biases, which sharpens the reading that the main effects reflect portfolio reorganization rather than a general change in investor type. The pattern in the main table is intentionally described as descriptive. Lower and moderate RA shares are associated with broader and less heavily traded self-directed sleeves. Very high RA shares are associated with a smaller residual self-directed sleeve that can look more concentrated or more actively traded. This non-monotonicity is consistent with the main paper's distinction between entry and intensity. A higher RA share does not simply scale up the same portfolio adjustment. Once a large share of fund wealth is delegated, the remaining self-directed sleeve is a selected residual portfolio, and its composition need not move monotonically with RA intensity.

A.11. Stability, Age Correlations, and Self-Directed Sleeve

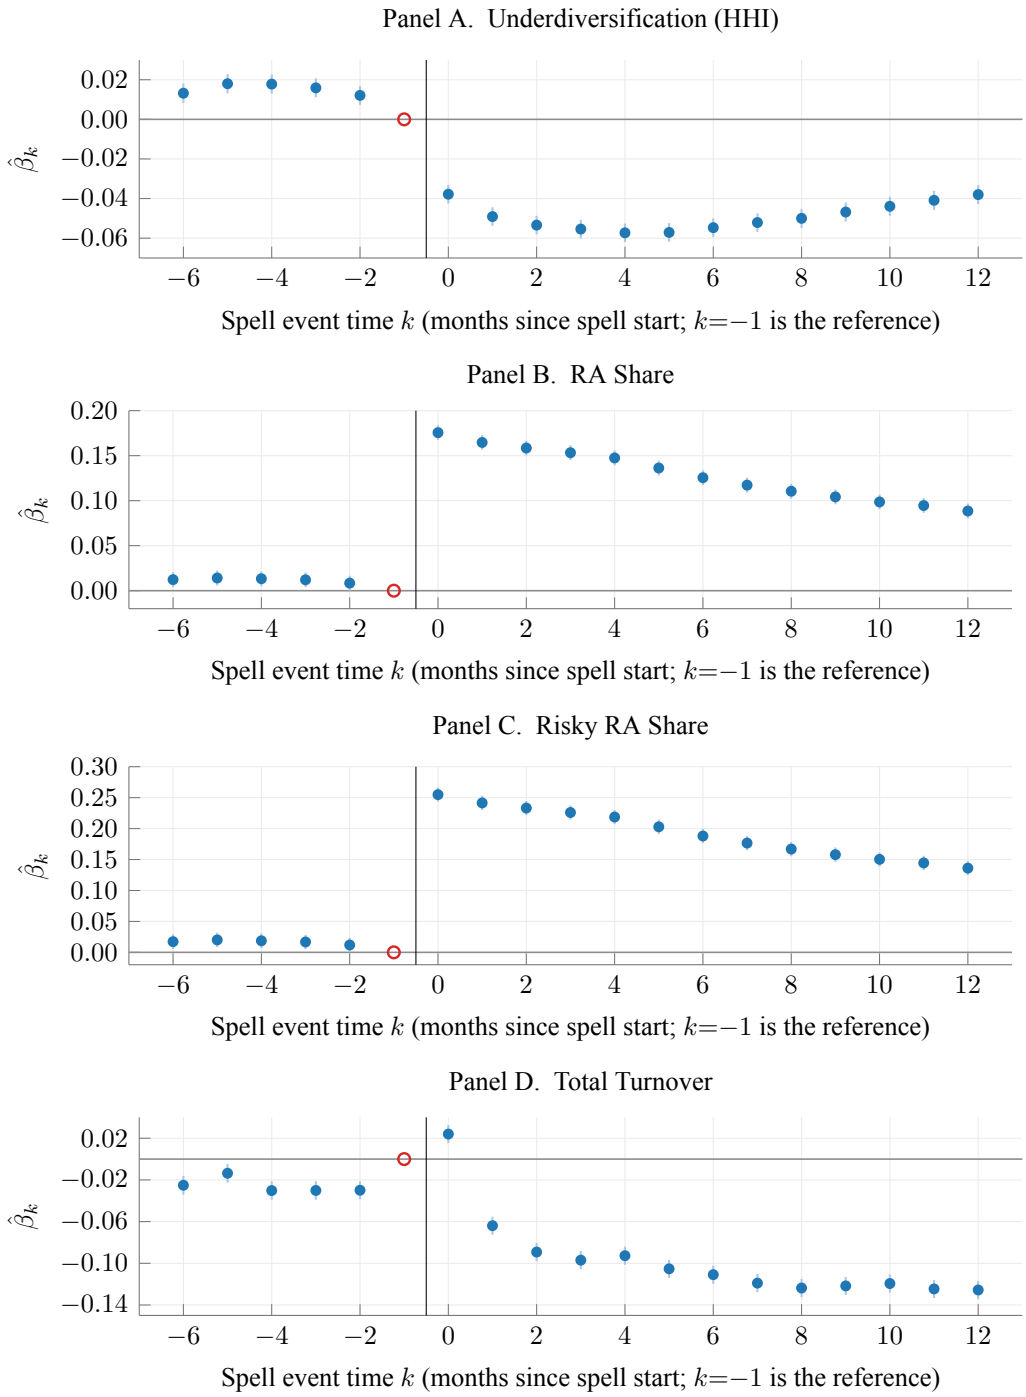
Online Appendix Table A.16 and Figure A.3 report stability checks across early and late sample windows and adoption cohorts. These exercises are scope checks; the main monthly Fama–MacBeth estimates remain the evidence. The early months of the sample correspond to a newly available product and a more selected adopter base, so some demographic coefficients are less stable in short windows; the early-period and late-cohort columns average only six and twenty-three monthly cross-sections respectively. Signs of the demographic effects are mostly, but not uniformly, stable. In the RA-share panel, log total financial assets is negative on the conditional share in all four splits, age is negative in all four, and cash-management assets are positive in all four; the female coefficient is negative everywhere except the early period, and risk preference moves from -0.019 early to $+0.008$ late, both significant. The risky-RA-share panel is less stable still: female moves from $+0.050$ to -0.002, risk preference from -0.059 to $+0.005$, and age from an insignificant $+0.006$ to -0.007, all between the early and late windows. The table is best read as a scope check on the wealth and age gradients; it does not support a stability claim for the demographic block as a whole. The magnitude of the wealth effect shrinks over the sample, from -0.216 in the early period to -0.088 in the late period, which is consistent with the wealth gradient on intensive delegation weakening as the technology matures. The main sign patterns are evaluated in the longer monthly panels and in the joint specification, not in six-month subsamples alone. Online Appendix Table A.17 reports correlations between age and the behavioral-bias, skill, and experience measures. These correlations support the interpretation of the age coefficient in the main text. Older investors trade less, with a panel correlation of -0.068 between age and total turnover; are less concentrated by HHI, at -0.034; and display stronger confirmation bias, at $+0.060$. They also hold more funds and have longer fund experience. Once these behavioral correlates are controlled for, the residual age coefficient operates mainly on adoption rather than on conditional delegation. The appendix table is therefore not a separate mechanism test; it explains why the raw age coefficient changes when the behavioral block enters. Online Appendix Table A.18 reports self-directed-sleeve composition within the RA-adopter sample. The lagged combined-advisor indicator, which captures use of delegated advisory products other than the focal BangNiTou RA product, is associated with more self-directed funds, lower self-directed concentration, a higher self-directed share, and lower total and speculative turnover. This evidence is descriptive and should not

be read as a causal effect of advisory use on self-directed discipline. It supports a more limited point: active delegated-product use can coexist with substantial and disciplined self-directed investing, so delegation does not mechanically replace the self-directed sleeve.

A.12. Appendix Figures

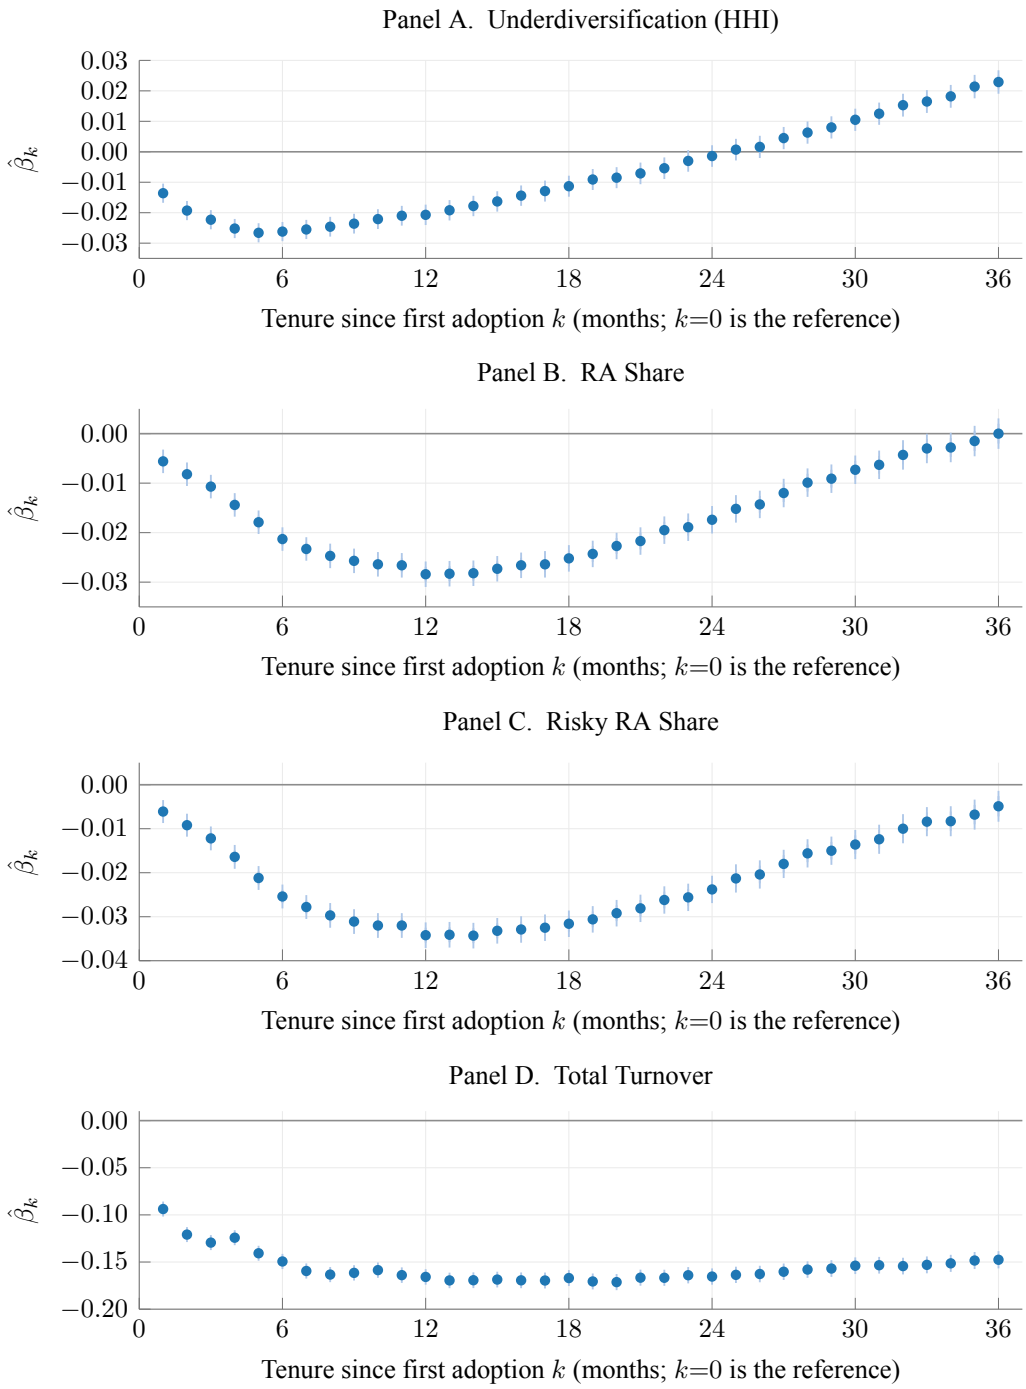
Figure A.1 repeats the event-study analysis for all adoption spells, including re-entry spells, instead of only first-ever adoption. Figure A.2 reports the post-adoption tenure event study for allocation and self-directed-sleeve outcomes. Figure A.3 plots time-stability and cohort-split coefficients. These figures are included to show that the main visual patterns are not driven by a single event definition, sample window, or return specification.

Figure A.1: Event Study for All Adoption Spells (Not Just the First)



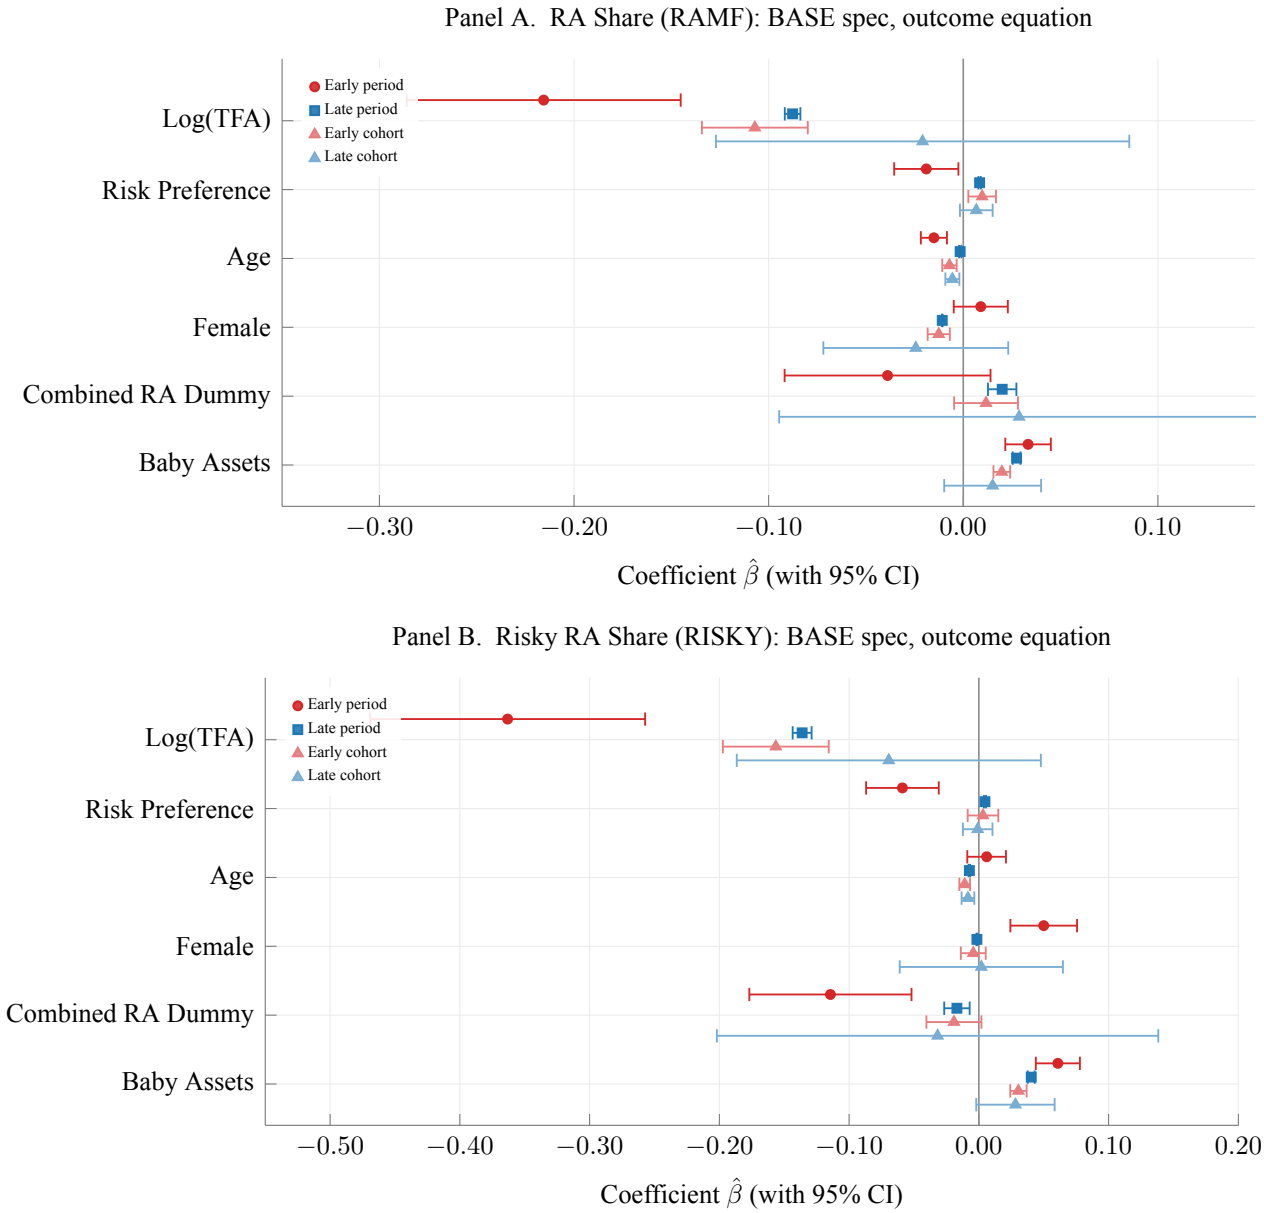
Notes. Event-study coefficients from **spell** fixed-effects regressions $y_{it} - \bar{y}_s = \sum_k \beta_k \cdot \mathbf{1}\{k_{st} = k\} + \varepsilon_{it}$, where s indexes RA-adoption *spells* (every $0 \to 1$ transition of RA_{it} opens a new spell). Demeaning is done within spell rather than within investor, so investors who exit and re-enter RA contribute multiple observations. Window is shorter than the first-adoption event study in Figure 4 ($k \in [-6, +12]$ vs. $[-12, +24]$) because spells can be re-entered; the relevant comparison is short-horizon. The spell sample is capped at 100,000 randomly drawn spells (seed 42) for computational tractability. Points are $\hat{\beta}_k$; vertical bars are 95% heteroskedasticity-robust CIs; the hollow red marker at $k = -1$ marks the reference period. Cross-panel consistency with the first-adoption event study (Figure 4) in the overlap $k \in [-6, +12]$ confirms that the within-investor first-adoption pattern is not driven by first-adoption effects alone.

Figure A.2: Post-Adoption Tenure Event Study



Notes. Event-study coefficients from within-investor fixed-effects regressions $y_{it} - \bar{y}_i = \sum_k \beta_k \cdot \mathbf{1}\{\text{tenure}_{it} = k\} + \varepsilon_{it}$, where tenure is months elapsed since an investor's first RA adoption. Sample restricted to the post-adoption window of first adopters; the reference period is $k = 0$ (month of first adoption). In contrast to Figure 4, which plots absolute event-time coefficients and shows the dominant level shift at adoption, this figure traces how the post-adoption gap from the adoption baseline *evolves with tenure*. Points are $\hat{\beta}_k$; vertical bars are 95% heteroskedasticity-robust CIs. All four outcomes drift back toward the adoption baseline as tenure grows: HHI recovers from -0.027 at month 5 to $+0.023$ at month 36 (a clear reconcentration); RA share recovers from -0.028 to ≈ 0; risky RA share from -0.034 to -0.005; turnover partially recovers from -0.170 to -0.148. The pattern is the within-adopter analogue of the "reclaiming control" result in Figure 3.

Figure A.3: Time-Varying Stability and Cohort-Split Coefficient Plot



Notes. BASE-spec Heckman outcome-equation coefficients $\hat{\beta}_k$ across four sample splits: (i) *Early period* — the first 6 months of the panel; (ii) *Late period* — the final 6 months; (iii) *Early cohort* — investors whose first RA adoption occurred in the first 6 months; (iv) *Late cohort* — investors whose first RA adoption occurred in the final 6 months. Panel A uses RAMF $=$ RA$/$(RA $+$ MF) as the outcome; Panel B uses the risky RA share. Points are $\hat{\beta}$ and horizontal bars are 95% CIs. The BASE spec includes only demographics (Log(TFA), Risk Preference, Age, Female, Combined RA Dummy, Baby Assets) — the most comparable subset across the four splits. Wide CIs in the Early-period and Late-cohort bins reflect small samples (only one cohort observed for less than 6 months), but the *sign pattern* of the demographic effects is stable across splits: Log(TFA) is consistently negative, Baby Assets consistently positive, and the Combined RA Dummy consistently positive in the late period and late cohort (consistent with adoption persistence).

A.13. Appendix Tables

Table A.1: Main Heckman Results: Weighted Estimation

	Panel 1: Baseline (T=48)		Panel 2: +Skills (T=47)		Panel 3: +Biases (T=36)	
Variable	Probit ME	Outcome (RA_{share})	Probit ME	Outcome (RA_{share})	Probit ME	Outcome (RA_{share})
Constant	-0.063***	0.756***	-0.063***	0.639**	-0.063***	0.458***
	(-184.010)	(2.715)	(-143.929)	(2.480)	(-333.818)	(10.688)
Log(Total Financial Assets)$_{t-1}$	0.006***	-0.175***	0.006***	-0.128***	0.003***	-0.058***
	(10.722)	(-4.389)	(10.938)	(-5.478)	(5.935)	(-10.843)
Risk Preference	0.001***	-0.012	0.001***	0.012*	0.001***	0.003***
	(5.548)	(-0.969)	(4.474)	(1.792)	(7.218)	(4.707)
Age	0.000	-0.004	0.000	-0.019***	0.001***	-0.011***
	(0.333)	(-0.538)	(0.284)	(-3.061)	(15.540)	(-14.330)
Female	-0.002***	0.014	-0.002***	0.002	-0.002***	0.010***
	(-12.853)	(1.229)	(-15.712)	(0.227)	(-18.635)	(8.898)
Combined RA Dummy$_{t-1}$	0.011***	-0.088***	0.010***	-0.095***	0.010***	-0.037***
	(25.807)	(-4.006)	(28.284)	(-3.626)	(67.314)	(-6.831)
Baby Assets$_{t-1}$	-0.002***	0.037***	-0.001***	0.030***	-0.001***	0.012***
	(-16.439)	(5.403)	(-15.119)	(6.552)	(-7.976)	(8.476)
Fund Experience (Months)$_{t-1}$			0.002***	-0.056***		
			(9.752)	(-10.974)		
Selection Ability$_{t-1}$			0.000*	-0.004***		
			(1.770)	(-2.888)		
Timing Ability$_{t-1}$			0.000	-0.002		
			(0.819)	(-1.566)		
Anchoring Bias$_{t-1}$					0.001***	-0.010***
					(11.006)	(-7.365)
Total Turnover$_{t-1}$					-0.001***	0.062***
					(-13.757)	(7.196)
Speculative Turnover$_{t-1}$					0.000***	-0.018***
					(-3.699)	(-8.071)
Attention Bias$_{t-1}$					0.000***	-0.001
					(-6.242)	(-1.637)
Disposition Effect$_{t-1}$					0.000***	-0.003***
					(5.265)	(-7.054)
Underdiversification HHI$_{t-1}$					-0.003***	0.055***
					(-98.504)	(26.405)
Extrapolation Bias$_{t-1}$					0.000***	0.001
					(-3.139)	(0.998)
Gambling Bias$_{t-1}$					0.000***	0.002***
					(-5.926)	(3.253)
Confirmation Bias$_{t-1}$					0.001***	-0.039***
					(24.914)	(-8.100)
Herding Bias$_{t-1}$					0.000***	0.000
					(-7.486)	(0.833)
Months Since Registration$_{t-1}$	0.001***		0.000***		0.001***	
	(9.771)		(7.564)		(8.828)	
Inverse Mills Ratio (λ)		-0.197**		-0.164*		-0.098***
		(-2.054)		(-1.718)		(-6.119)
Province Dummies	Yes	Yes	Yes	Yes	Yes	Yes
(Pseudo) R^2	0.044	0.132	0.041	0.146	0.046	0.228
Adjusted R^2		0.131		0.145		0.227
$\hat{\rho}$ (mean / median)		-0.364 / -0.472		-0.702 / -0.844		-0.463 / -0.463
$\hat{\sigma}$ (mean / median)		0.348 / 0.243		0.386 / 0.352		0.203 / 0.200
$\hat{\lambda}$ (mean / median)		-0.197 / -0.113		-0.164 / -0.287		-0.098 / -0.097
N (monthly mean)	99,992	32,587	95,129	32,747	91,845	34,180
N (total)	4,799,616	1,564,155	4,471,044	1,539,119	3,306,408	1,230,489

Notes. This table reports a robustness check on the main (unweighted) results of Table 2, re-estimating the same specifications with calibrated inverse-probability weights. Estimates are from weighted Heckman two-stage (Tobit-2) selection models estimated monthly and averaged following Fama and MacBeth (1973). Both stages apply weights that align the sample RA-adoption rate with an assumed population rate of $p_{pop} = 1/101 \approx 0.99\%$: within each month, RA-adopter observations receive weight $p_{pop}/p_{samp,t}$ and non-adopter observations $(1-p_{pop})/(1-p_{samp,t})$, normalised to mean one. The selection stage is a weighted Probit and the outcome stage is WLS

on the adopter subsample. Reported selection-stage values are *approximate marginal effects* on the probability of RA adoption, computed as $\text{ME}_k = \phi(\Phi^{-1}(p_{pop})) \cdot \hat{\beta}_k$, where $\hat{\beta}_k$ is the Fama–MacBeth average of the monthly weighted Probit coefficients. Because the calibrated weights centre the fitted adoption probability on the population rate rather than the sample rate, the marginal effects are evaluated at $p_{pop} = 1/101$, giving a scaling factor $\phi(\Phi^{-1}(1/101)) = 0.026$. Marginal effects are computed from the unrounded coefficients and scaling factor and then rounded to three decimals for display; because the scaling factor is small, several marginal effects round to 0.000 even where the underlying Probit coefficient is precisely estimated, and significance should be read from the t-statistics rather than from the displayed magnitudes. These marginal effects are evaluated at a population base rate of roughly one percent and are therefore not directly comparable in magnitude to those in Table 2, which are evaluated at the sample adoption rate of the estimation panel. Outcome-stage coefficients are reported as estimated (no scaling). Newey–West HAC t-statistics (max lag = 3), computed from the time series of monthly coefficients, are in parentheses; t-statistics, stars, and significance are invariant to the scaling. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed). Panel 1 reports the baseline specification. Panel 2 adds investor skills (fund experience, selection ability, timing ability). Panel 3 adds behavioral biases. The selection (Probit) dependent variable is $\text{RA}_{it} = 1$ if the investor holds robo-advisor assets; the outcome dependent variable is RA_{share} = RA assets / (RA + mutual fund assets), observed only for RA-adopter months. All right-hand-side variables except Age, Female, and Risk Preference are lagged one period. Months Since Registration enters the selection equation only and serves as the exclusion restriction; the Inverse Mills Ratio, constructed from the monthly weighted Probit, enters the outcome equation only. Province dummies are included in both equations but not reported. The (Pseudo) R^2 row reports the McFadden pseudo-R^2 in the selection columns and the OLS R^2 in the outcome columns, averaged across months. T is the number of monthly cross-sections with successful estimation in each specification. N (total) is the sum of monthly observations across the T months; N (monthly mean) is the corresponding per-month average; reported N are unweighted observation counts. $\hat{\rho}$ is the average estimated correlation between the selection and outcome errors, $\hat{\sigma}$ the average standard deviation of the outcome error, and $\hat{\lambda} = \hat{\rho}\hat{\sigma}$ the average coefficient on the Inverse Mills Ratio; means and medians are taken over the T monthly estimates.

Table A.2: Exclusion-Restriction Sensitivity: Alternative Placements of Platform Tenure

	Arm B: Tenure in both equations (T=36)						Arm C: Tenure in neither equation (T=36)						Arm E: Benchmark—tenure in selection only (T=36)					
	Baseline		+Skills		+Biases		Baseline		+Skills		+Biases		Baseline		+Skills		+Biases	
Variable	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.
Constant	-0.408***	-0.637***	-0.416***	-0.542***	-0.401***	-0.605***	-0.402***	-0.656***	-0.412***	-0.569***	-0.395***	-0.618***	-0.408***	-0.002	-0.416***	0.001	-0.401***	-0.073
	(-13.834)	(-9.974)	(-13.476)	(-7.290)	(-12.728)	(-9.799)	(-13.445)	(-9.942)	(-13.163)	(-7.309)	(-12.377)	(-9.925)	(-13.834)	(-0.036)	(-13.476)	(0.014)	(-12.728)	(-1.578)
Log(Total Financial Assets)$_{t-1}$	0.358***	0.091***	0.342***	0.070***	0.209***	0.063***	0.363***	0.102***	0.346***	0.082***	0.215***	0.070***	0.358***	-0.055***	0.342***	-0.052***	0.209***	-0.008
	(9.957)	(4.496)	(10.221)	(3.315)	(6.140)	(3.326)	(10.144)	(4.645)	(10.367)	(3.493)	(6.312)	(3.480)	(9.957)	(-6.615)	(10.221)	(-5.273)	(6.140)	(-1.279)
Risk Preference	0.054***	0.030***	0.038***	0.024***	0.050***	0.028***	0.060***	0.034***	0.041***	0.027***	0.056***	0.031***	0.054***	0.012***	0.038***	0.015***	0.050***	0.014***
	(6.885)	(11.854)	(4.737)	(15.392)	(6.611)	(10.516)	(7.890)	(14.884)	(5.403)	(18.344)	(7.627)	(13.343)	(6.885)	(7.238)	(4.737)	(8.028)	(6.611)	(9.872)
Age	0.043***	0.014***	0.034***	0.010***	0.044***	0.012***	0.048***	0.018***	0.039***	0.013***	0.049***	0.016***	0.043***	-0.002	0.034***	0.000	0.044***	0.001
	(15.922)	(4.140)	(13.468)	(2.921)	(19.385)	(3.953)	(16.508)	(4.844)	(14.155)	(3.488)	(19.730)	(4.631)	(15.922)	(-0.833)	(13.468)	(0.151)	(19.385)	(0.416)
Female	-0.104***	-0.058***	-0.107***	-0.052***	-0.138***	-0.064***	-0.098***	-0.056***	-0.102***	-0.051***	-0.133***	-0.062***	-0.104***	-0.016***	-0.107***	-0.016***	-0.138***	-0.018***
	(-20.358)	(-11.162)	(-22.027)	(-9.603)	(-19.661)	(-9.820)	(-17.933)	(-10.579)	(-19.689)	(-9.326)	(-17.841)	(-9.504)	(-20.358)	(-4.919)	(-22.027)	(-4.188)	(-19.661)	(-4.243)
Combined RA Dummy$_{t-1}$	0.660***	0.280***	0.633***	0.239***	0.576***	0.251***	0.663***	0.290***	0.634***	0.250***	0.578***	0.257***	0.660***	0.044**	0.633***	0.048**	0.576***	0.084***
	(77.541)	(10.692)	(74.834)	(8.280)	(90.326)	(11.166)	(76.757)	(10.657)	(74.400)	(8.245)	(89.735)	(11.236)	(77.541)	(2.502)	(74.834)	(2.379)	(90.326)	(5.318)
Baby Assets$_{t-1}$	-0.094***	-0.027***	-0.089***	-0.020***	-0.054***	-0.018***	-0.094***	-0.028***	-0.089***	-0.022***	-0.054***	-0.019***	-0.094***	0.011***	-0.089***	0.010**	-0.054***	-0.001
	(-14.191)	(-3.910)	(-12.379)	(-3.074)	(-9.797)	(-3.357)	(-14.110)	(-4.075)	(-12.268)	(-3.271)	(-9.715)	(-3.442)	(-14.191)	(2.852)	(-12.379)	(2.370)	(-9.797)	(-0.151)
Fund Experience (Months)$_{t-1}$			0.108***	0.017***					0.111***	0.022***					0.108***	-0.022***		
			(20.231)	(4.297)					(20.195)	(5.924)					(20.231)	(-3.653)		
Selection Ability$_{t-1}$			0.018***	0.008***					0.019***	0.008***					0.018***	0.002		
			(3.603)	(2.689)					(3.598)	(2.743)					(3.603)	(1.042)		
Timing Ability$_{t-1}$			0.004	0.000					0.004	0.000					0.004	-0.002		
			(1.012)	(0.058)					(1.003)	(0.113)					(1.012)	(-1.595)		
Anchoring Bias$_{t-1}$					0.056***	0.022***					0.056***	0.022***					0.056***	0.003*
					(12.206)	(7.242)					(12.027)	(7.252)					(12.206)	(1.648)
Total Turnover$_{t-1}$					-0.083***	0.009					-0.084***	0.007					-0.083***	0.039***
					(-14.562)	(1.125)					(-14.479)	(0.944)					(-14.562)	(5.475)
Attention Bias$_{t-1}$					-0.023***	-0.013***					-0.023***	-0.014***					-0.023***	-0.006***
					(-7.107)	(-6.631)					(-7.275)	(-6.785)					(-7.107)	(-4.689)
Disposition Effect$_{t-1}$					0.007***	0.001					0.008***	0.001					0.007***	-0.002***
					(5.256)	(0.859)					(5.293)	(1.090)					(5.256)	(-4.422)
Underdiversification HHI$_{t-1}$					-0.172***	-0.042***					-0.170***	-0.043***					-0.172***	0.018***
					(-95.296)	(-5.121)					(-91.973)	(-5.177)					(-95.296)	(3.153)
Extrapolation Bias$_{t-1}$					-0.013***	-0.006**					-0.013***	-0.006**					-0.013***	-0.002*
					(-2.859)	(-2.493)					(-2.904)	(-2.548)					(-2.859)	(-1.755)
Gambling Bias$_{t-1}$					-0.024***	-0.013***					-0.024***	-0.013***					-0.024***	-0.004***
					(-5.861)	(-5.043)					(-5.862)	(-5.057)					(-5.861)	(-3.471)
Confirmation Bias$_{t-1}$					0.080***	0.006					0.081***	0.007*					0.080***	-0.023***
					(25.986)	(1.339)					(25.842)	(1.692)					(25.986)	(-5.717)
Herding Bias$_{t-1}$					-0.017***	-0.009***					-0.017***	-0.009***					-0.017***	-0.004***
					(-7.484)	(-5.254)					(-7.575)	(-5.239)					(-7.484)	(-3.766)
Months Since Registration$_{t-1}$	0.038***	0.023***	0.030***	0.019***	0.036***	0.021***							0.038***		0.030***		0.036***	
	(12.403)	(9.932)	(10.383)	(7.295)	(10.968)	(8.933)							(12.403)		(10.383)		(10.968)	
Inverse Mills Ratio (λ)		0.769***		0.684***		0.750***		0.791***		0.713***		0.765***		0.180***		0.182***		0.255***
		(10.829)		(8.312)		(11.009)		(10.654)		(8.177)		(10.965)		(3.596)		(3.193)		(5.403)
Province Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
(Pseudo) R^2	0.070	0.110	0.074	0.127	0.090	0.234	0.069	0.109	0.074	0.127	0.089	0.234	0.070	0.105	0.074	0.123	0.090	0.230
Adjusted R^2		0.109		0.126		0.233		0.108		0.126		0.233		0.104		0.122		0.229
$\hat{\rho}$		3.889		3.452		4.116		3.991		3.590		4.196		0.860		0.875		1.381
$\hat{\sigma}$		0.195		0.198		0.180		0.195		0.193		0.180		0.196		0.194		0.181
$\hat{\lambda}$ (IMR coefficient)		0.769		0.684		0.750		0.791		0.713		0.765		0.180		0.182		0.255
IMR-on-X R^2		0.995		0.995		0.995		0.995		0.995		0.995		0.985		0.990		0.988

Continued on next page

Table A.2 (continued)

	Arm B: Tenure in both equations (T=36)						Arm C: Tenure in neither equation (T=36)						Arm E: Benchmark—tenure in selection only (T=36)					
	Baseline		+Skills		+Biases		Baseline		+Skills		+Biases		Baseline		+Skills		+Biases	
Variable	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.	Sel.	Out.
Condition number		141.5		143.3		137.7		142.2		143.1		137.7		90.1		106.9		96.1
N (monthly mean)	91,813	34,173	91,813	34,173	91,813	34,173	91,813	34,173	91,813	34,173	91,813	34,173	91,813	34,173	91,813	34,173	91,813	34,173
N (total)	3,305,265	1,230,240	3,305,265	1,230,240	3,305,265	1,230,240	3,305,265	1,230,240	3,305,265	1,230,240	3,305,265	1,230,240	3,305,265	1,230,240	3,305,265	1,230,240	3,305,265	1,230,240

Notes. This table assesses the sensitivity of the Heckman two-step estimates to the placement of the exclusion restriction (Months Since Registration, i.e. platform tenure). Arm B includes tenure in both the selection and outcome equations; Arm C includes it in neither; Arm E is the benchmark with tenure in the selection equation only, corresponding to the specification of Table 2. For each specification, Sel. reports the selection equation (monthly Probit, raw coefficients, not marginal effects and therefore not directly comparable in magnitude to the marginal effects of Table 2) and Out. reports the outcome equation (adopter-month OLS including the inverse Mills ratio). Coefficients are Fama–MacBeth means of monthly estimates, reported to three decimals; Newey–West t-statistics with lag $\min(3, T-1)$, computed from the time series of monthly coefficients, are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed). Because Arms B and E share an identical selection equation, their Sel. columns coincide by construction; Arm C differs only through the omission of tenure. The sample is identical across the three arms, estimation is unweighted, and all right-hand-side variables except Age, Female, and Risk Preference are lagged one month; province dummies enter both equations but are not reported. Arm C is identified only through the Probit functional-form assumptions (the nonlinearity of the inverse Mills ratio), with no excluded instrument. The (Pseudo) R^2 row reports the McFadden pseudo-R^2 in Sel. columns and the OLS R^2 in Out. columns, averaged across months. $\hat{\sigma}$ is the average standard deviation of the outcome error, $\hat{\lambda}$ the average coefficient on the inverse Mills ratio, and $\hat{\rho}$ the average of the monthly ratios $\hat{\lambda}_t/\hat{\sigma}_t$. Values of $\hat{\rho}$ far outside the admissible correlation range $[-1, 1]$ in Arms B and C indicate that the selection correction is not separately identified when the exclusion restriction is absent from (or duplicated in) the outcome equation. Consistent with this, the IMR-on-X R^2 — the R^2 from regressing the inverse Mills ratio on the outcome regressors — is near one in all arms and highest in B and C, and the condition number of the outcome design matrix rises relative to the benchmark, confirming that identification in Arms B and C rests on residual nonlinearity alone. In Arm E, the Months Since Registration selection coefficient is the exclusion-relevance statistic and is strongly significant in all specifications. N (total) is the sum of monthly observations across the T months; N (monthly mean) is the corresponding per-month average.

Table A.3: Heckman Two-Stage Estimates: Full Specification (Baseline + Skills + Biases)

Variable	Probit ME	Outcome (RA_{share})
Constant	-0.152***	-0.095*
	(-12.461)	(-1.870)
Log(Total Financial Assets)$_{t-1}$	0.079***	-0.005
	(6.081)	(-0.664)
Risk Preference	0.015***	0.015***
	(5.587)	(12.132)
Age	0.015***	0.000
	(21.152)	(0.011)
Female	-0.052***	-0.021***
	(-19.747)	(-4.393)
Combined RA Dummy$_{t-1}$	0.214***	0.090***
	(93.140)	(5.316)
Baby Assets$_{t-1}$	-0.020***	-0.002
	(-9.001)	(-0.460)
Fund Experience (Months)$_{t-1}$	0.024***	-0.005
	(10.302)	(-1.413)
Selection Ability$_{t-1}$	0.007***	0.003*
	(3.412)	(1.896)
Timing Ability$_{t-1}$	0.001	-0.001
	(0.867)	(-0.823)
Anchoring Bias$_{t-1}$	0.022***	0.003
	(12.667)	(1.362)
Total Turnover$_{t-1}$	-0.028***	0.039***
	(-12.421)	(6.129)
Attention Bias$_{t-1}$	-0.008***	-0.006***
	(-6.684)	(-4.577)
Disposition Effect$_{t-1}$	0.001**	-0.000
	(2.275)	(-0.719)
Underdiversification HHI$_{t-1}$	-0.065***	0.016***
	(-91.321)	(2.610)
Extrapolation Bias$_{t-1}$	-0.004***	-0.002*
	(-3.075)	(-1.707)
Gambling Bias$_{t-1}$	-0.009***	-0.004***
	(-5.798)	(-2.820)
Confirmation Bias$_{t-1}$	0.019***	-0.011***
	(21.864)	(-3.264)
Herding Bias$_{t-1}$	-0.006***	-0.004***
	(-6.519)	(-3.807)
Months Since Registration$_{t-1}$	0.012***	
	(10.176)	
Inverse Mills Ratio (λ)		0.277***
		(5.404)
Province Dummies	Yes	Yes
Pseudo R^2 / R^2	0.092	0.226
N	3,214,651	1,199,341

Notes. Estimates are from an unweighted Heckman two-stage (Tobit-2) selection model estimated monthly. Reported selection-stage values are *approximate marginal effects* on the probability of RA adoption, computed as $\text{ME}_k = \phi(\Phi^{-1}(\bar{P})) \cdot \hat{\beta}_k$, where $\bar{P} = 0.373$ is the average monthly RA-active rate in the sample and $\hat{\beta}_k$ is the Fama–MacBeth average Probit coefficient. The scaling factor is $\phi(\Phi^{-1}(0.373)) = 0.379$. Outcome-stage coefficients are reported as estimated (no scaling). Newey–West HAC t-statistics (max lag = 3) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed). The full specification includes baseline demographics and assets, investor skills (fund experience, selection ability, timing ability), and behavioral biases (anchoring, turnover, attention, disposition effect, underdiversification, extrapolation, gambling, confirmation, herding). The selection (Probit) dependent variable is $\text{RA}_{it} = 1$ if the investor holds robo-advisor assets; the outcome dependent variable is RA_{share} = RA assets / (RA + mutual fund assets). Months Since Registration serves as the exclusion restriction (selection-only). All right-hand-side variables except Age, Female, and Risk Preference are lagged one period. Province dummies are included but not reported. $\hat{\sigma} = 0.182$, $\hat{\lambda} = 0.277$ (FM averages). The two-step estimator does not constrain the implied $\hat{\rho} = \hat{\lambda}/\hat{\sigma}$ to $[-1, 1]$; its average exceeds one, reflecting the heavily loaded inverse-Mills term rather than a literal error correlation, so we report $\hat{\lambda}$ as the selection statistic. N for the outcome equation is the number of RA-adopter months; N for the selection equation is the full sample of investor-months. The $\phi(\Phi^{-1}(\cdot))$ approximation is accurate to within 1–5% of a properly computed average marginal effect for Probit models with well-behaved $X\hat{\beta}$ dispersion.

Table A.4: Alternative Outcome Definitions for the Intensive Margin

Panel A: DV-A $= \log(1 +$ **RA assets**$)$ (level-form intensive margin)						
	BASE		SKILL		BIAS	
Variable	Selection ME	Outcome	Selection ME	Outcome	Selection ME	Outcome
Constant	-0.197***	6.391***	-0.190***	7.424***	-0.150***	8.099***
	(-6.745)	(5.278)	(-7.655)	(10.299)	(-12.626)	(17.804)
Log(TFA) (t-1)	0.141***	1.720***	0.134***	1.566***	0.080***	1.591***
	(12.559)	(6.428)	(11.754)	(9.478)	(6.036)	(19.073)
Risk Preference	0.029***	0.221**	0.022***	0.166***	0.019***	0.092***
	(5.616)	(2.308)	(4.261)	(3.666)	(6.626)	(5.682)
Age	0.004	-0.200***	0.003	-0.157***	0.017***	-0.137***
	(0.770)	(-3.284)	(0.612)	(-4.086)	(25.518)	(-5.148)
Female	-0.042***	-0.003	-0.043***	0.080	-0.052***	0.182***
	(-17.144)	(-0.026)	(-19.983)	(1.324)	(-19.617)	(8.010)
Combined RA Dummy (t-1)	0.224***	-0.292	0.222***	-0.471***	0.219***	-0.554***
	(24.914)	(-1.312)	(29.157)	(-3.061)	(98.397)	(-4.039)
Baby Assets (t-1)	-0.037***	-0.077	-0.034***	-0.038	-0.020***	-0.048***
	(-18.857)	(-1.408)	(-16.497)	(-1.235)	(-9.378)	(-2.607)
Fund Experience (mos)			0.035***	-0.216***		
			(9.546)	(-10.657)		
Selection Skill			0.004*	-0.011		
			(1.886)	(-1.432)		
Timing Rate			0.001	-0.005		
			(0.968)	(-0.940)		
Anchoring Bias					0.022***	-0.052***
					(12.976)	(-3.618)
Total Turnover					-0.031***	0.069***
					(-14.565)	(4.835)
Attention Bias					-0.009***	0.015
					(-6.870)	(1.456)
Disposition Effect					0.001*	0.013***
					(1.751)	(3.508)
Underdiversification (HHI)					-0.066***	0.204***
					(-99.536)	(4.328)
Extrapolation Bias					-0.005***	0.011*
					(-2.731)	(1.869)
Gambling Bias					-0.009***	-0.010*
					(-6.546)	(-1.769)
Confirmation Bias					0.022***	-0.048***
					(29.161)	(-3.161)
Herding Bias					-0.006***	0.029***
					(-7.322)	(4.698)
Months Since Reg. (t-1)	0.012***		0.008***		0.014***	
	(10.538)		(7.800)		(11.047)	
Inverse Mills Ratio		-0.155		-0.838		-1.411***
		(-0.196)		(-1.623)		(-3.894)
Province Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2 / R^2	—	0.251	—	0.258	—	0.250
N	4,799,616	1,564,155	4,471,044	1,539,119	3,214,908	1,199,373

Notes to Panel A. Outcome variable is $\log(1 +$ ra_assets_raw$)$, the log of (one plus) the level of RA assets. This preserves variation in absolute RA holdings that is normalised out by the RA-share outcome in the main paper. Three nested specifications: BASE (demographics), SKILL (adds Fund Experience, Selection Skill, Timing Rate), and BIAS (adds 9 biases). Heckman two-stage with Fama–MacBeth monthly estimation and Newey–West HAC t-statistics (max lag $= 3$) in parentheses. The outcome is estimated only on the sample with RA $= 1$. Selection-stage values are approximate marginal effects on the probability of RA adoption, $\text{ME}_k = \phi(\Phi^{-1}(\bar{P}))\hat{\beta}_k$, with $\bar{P} = 0.326, 0.344, 0.373$ (factors $0.360, 0.368, 0.379$) for BASE/SKILL/BIAS on the 4.9-million investor-month panel; outcome coefficients are unscaled and t-statistics and significance are invariant. $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.10$.

Panel B: DV-B = RA / Total Financial Assets (share-of-wealth intensive margin)

	BASE		SKILL		BIAS	
Variable	Selection ME	Outcome	Selection ME	Outcome	Selection ME	Outcome
Constant	-0.197***	0.000	-0.190***	-0.045	-0.150***	0.056*
	(-6.745)	(0.000)	(-7.655)	(-0.759)	(-12.626)	(1.758)
Log(TFA) (t-1)	0.141***	-0.025***	0.134***	-0.010	0.080***	-0.011***
	(12.559)	(-4.533)	(11.754)	(-0.812)	(6.036)	(-2.725)
Risk Preference	0.029***	0.014***	0.022***	0.018***	0.019***	0.008***
	(5.616)	(5.600)	(4.261)	(3.997)	(6.626)	(8.215)
Age	0.004	-0.008***	0.003	-0.010**	0.017***	-0.004***
	(0.770)	(-2.908)	(0.612)	(-2.064)	(25.518)	(-2.780)
Female	-0.042***	-0.012***	-0.043***	-0.015***	-0.052***	-0.004*
	(-17.144)	(-4.005)	(-19.983)	(-2.832)	(-19.617)	(-1.737)
Combined RA Dummy (t-1)	0.224***	0.014	0.222***	0.024*	0.219***	0.010
	(24.914)	(1.429)	(29.157)	(1.952)	(98.397)	(0.883)
Baby Assets (t-1)	-0.037***	-0.010***	-0.034***	-0.011***	-0.020***	-0.010***
	(-18.857)	(-4.385)	(-16.497)	(-4.062)	(-9.378)	(-6.239)
Fund Experience (mos)			0.035***	-0.013***		
			(9.546)	(-5.581)		
Selection Skill			0.004*	0.000		
			(1.886)	(-0.405)		
Timing Rate			0.001	-0.001		
			(0.968)	(-1.435)		
Anchoring Bias					0.022***	0.000
					(12.976)	(-0.013)
Total Turnover					-0.031***	0.017***
					(-14.565)	(7.329)
Attention Bias					-0.009***	-0.002***
					(-6.870)	(-2.581)
Disposition Effect					0.001*	0.002***
					(1.751)	(5.882)
Underdiversification (HHI)					-0.066***	0.017***
					(-99.536)	(4.912)
Extrapolation Bias					-0.005***	0.000
					(-2.731)	(0.356)
Gambling Bias					-0.009***	-0.002***
					(-6.546)	(-2.621)
Confirmation Bias					0.022***	-0.006***
					(29.161)	(-4.410)
Herding Bias					-0.006***	-0.001
					(-7.322)	(-1.631)
Months Since Reg. (t-1)	0.012***		0.008***		0.014***	
	(10.538)		(7.800)		(11.047)	
Inverse Mills Ratio		0.110***		0.136***		0.065*
		(3.838)		(3.209)		(1.951)
Province Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2 / R^2	—	0.088	—	0.092	—	0.107
N	4,799,616	1,564,155	4,471,044	1,539,119	3,214,908	1,199,373

Notes to Panel B. Outcome variable is RA assets divided by total financial assets (RA/TFA) — the share-of-wealth intensive margin (as opposed to the main paper's RA/(RA+MF), which is the within-risky-portfolio share). Heckman two-stage with Fama–MacBeth monthly estimation and Newey–West HAC t-statistics (max lag = 3) in parentheses. $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.10$. Relative to the main-paper RA/(RA+MF) outcome (Table 2), the RA/TFA formulation shows substantially smaller coefficient magnitudes because it dilutes variation over all financial assets rather than just the risky sleeve, but the sign patterns on Risk Preference (+), Age (−), Female (−), and Baby Assets (−) are preserved. The bias coefficients on Total Turnover and Underdiversification both enter with positive signs in the RA/TFA outcome, consistent with the main-paper finding that high-turnover and concentrated investors hold larger absolute RA positions.

Table A.5: Control-Function Panel Heckman Estimates of Robo-Advisor Allocation

	Baseline ($T = 48$)		+Skills ($T = 47$)		+Biases ($T = 35$)	
Variable	Entity	Entity+Time	Entity	Entity+Time	Entity	Entity+Time
Panel A: RA Share $= RA/(RA + MF)$						
Log(Total Financial Assets)$_{t-1}$	-0.162^{***}	-0.141^{***}	-0.151^{***}	-0.120^{***}	-0.144^{***}	-0.118^{***}
	(-25.235)	(-30.240)	(-31.405)	(-21.821)	(-43.568)	(-33.191)
Combined RA Dummy$_{t-1}$	-0.011^{***}	-0.000	-0.002	0.006	-0.017^{***}	-0.000
	(-2.906)	(-0.074)	(-0.580)	(1.440)	(-5.689)	(-0.085)
Baby Assets$_{t-1}$	0.018^{***}	0.016^{***}	0.017^{***}	0.014^{***}	0.012^{***}	0.010^{***}
	(23.442)	(22.398)	(26.690)	(19.443)	(24.477)	(19.864)
Fund Experience (Months)$_{t-1}$			0.010^{***}	-0.259^{***}		
			(5.415)	(-17.450)		
Selection Ability$_{t-1}$			0.000	-0.000		
			(0.416)	(-0.755)		
Timing Ability$_{t-1}$			-0.000	-0.000		
			(-0.092)	(-1.078)		
Anchoring Bias$_{t-1}$					-0.006^{***}	-0.006^{***}
					(-10.298)	(-10.855)
Total Turnover$_{t-1}$					0.040^{***}	0.041^{***}
					(25.156)	(26.169)
Attention Bias$_{t-1}$					-0.002^{***}	-0.001^{***}
					(-3.991)	(-3.238)
Disposition Effect$_{t-1}$					-0.001	-0.001^{*}
					(-1.215)	(-1.752)
Underdiversification HHI$_{t-1}$					0.045^{***}	0.032^{***}
					(27.657)	(17.379)
Extrapolation Bias$_{t-1}$					-0.001^{***}	-0.000
					(-2.762)	(-1.337)
Gambling Bias$_{t-1}$					0.002^{***}	0.002^{***}
					(4.341)	(4.691)
Confirmation Bias$_{t-1}$					-0.015^{***}	-0.014^{***}
					(-15.545)	(-14.724)
Herding Bias$_{t-1}$					-0.003^{***}	-0.002^{***}
					(-8.560)	(-4.317)
Inverse Mills Ratio (λ)	-0.014	0.036^{***}	0.020	0.064^{***}	-0.118^{***}	-0.020^{*}
	(-0.774)	(2.967)	(1.462)	(3.961)	(-19.540)	(-1.887)
Demographics (risk, age, female)			Absorbed by entity FE			
Entity (investor) FE	Yes	Yes	Yes	Yes	Yes	Yes
Time (month) FE	No	Yes	No	Yes	No	Yes
N (selected)	1,564,155	1,564,155	1,539,119	1,539,119	1,199,373	1,199,373
Panel B: Risky RA Share $= RA_{\text{risky}}/(RA_{\text{risky}} + MF_{\text{risky}})$						
Log(Total Financial Assets)$_{t-1}$	-0.178^{***}	-0.159^{***}	-0.171^{***}	-0.140^{***}	-0.160^{***}	-0.135^{***}
	(-24.818)	(-30.313)	(-31.635)	(-24.986)	(-44.668)	(-34.180)
Combined RA Dummy$_{t-1}$	-0.013^{***}	-0.004	-0.004	0.003	-0.021^{***}	-0.005
	(-2.698)	(-0.920)	(-1.077)	(0.662)	(-6.013)	(-1.507)
Baby Assets$_{t-1}$	0.021^{***}	0.019^{***}	0.020^{***}	0.017^{***}	0.014^{***}	0.012^{***}
	(22.709)	(22.315)	(26.760)	(21.006)	(25.618)	(21.984)
Fund Experience (Months)$_{t-1}$			0.009^{***}	-0.245^{***}		
			(4.236)	(-16.457)		
Selection Ability$_{t-1}$			0.000	-0.001		
			(0.314)	(-1.279)		
Timing Ability$_{t-1}$			-0.000	-0.001		
			(-0.054)	(-1.299)		
Anchoring Bias$_{t-1}$					-0.007^{***}	-0.007^{***}
					(-10.757)	(-11.302)
Total Turnover$_{t-1}$					0.045^{***}	0.046^{***}
					(25.143)	(26.188)
Attention Bias$_{t-1}$					-0.002^{***}	-0.001^{***}
					(-4.035)	(-3.464)
Disposition Effect$_{t-1}$					-0.000	-0.001
					(-0.684)	(-1.120)
Underdiversification HHI$_{t-1}$					0.052^{***}	0.038^{***}
					(29.067)	(18.500)

Continued on next page

Table A.5 (continued)

	Baseline ($T = 48$)		+Skills ($T = 47$)		+Biases ($T = 35$)	
Variable	Entity	Entity+Time	Entity	Entity+Time	Entity	Entity+Time
Extrapolation Bias$_{t-1}$					-0.001^{***}	-0.000
					(-2.583)	(-1.461)
Gambling Bias$_{t-1}$					0.002^{***}	0.002^{***}
					(4.377)	(5.031)
Confirmation Bias$_{t-1}$					-0.015^{***}	-0.015^{***}
					(-17.142)	(-16.464)
Herding Bias$_{t-1}$					-0.004^{***}	-0.002^{***}
					(-8.579)	(-4.922)
Inverse Mills Ratio (λ)	-0.017	0.024^{*}	0.012	0.055^{***}	-0.145^{***}	-0.048^{***}
	(-0.840)	(1.651)	(0.714)	(3.089)	(-21.095)	(-4.107)
Demographics (risk, age, female)			Absorbed by entity FE			
Entity (investor) FE	Yes	Yes	Yes	Yes	Yes	Yes
Time (month) FE	No	Yes	No	Yes	No	Yes
N (selected)	1,564,316	1,564,316	1,539,280	1,539,280	1,199,511	1,199,511

Notes. The table reports Wooldridge (1995) / Semykina and Wooldridge (2010) control-function panel selection estimates. The dependent variable in Panel A is the robo-advisor allocation share $ra_share_{i,t} = RA_{i,t}/(RA_{i,t}+MF_{i,t})$, observed on robo-advisor-active investor-months; in Panel B it is the risky allocation share $RA_{\text{risky}}/(RA_{\text{risky}}+MF_{\text{risky}})$, observed on months with positive risky robo-advisor holdings. The two dependent variables use different selection rules, so their selected samples differ slightly. Stage 1 is a per-calendar-month cross-sectional probit of the selection indicator on the specification regressors, months-since-registration (the exclusion restriction), and Mundlak per-investor time-averages of the time-varying regressors. The inverse Mills ratio $\widehat{\lambda}_{i,t} = \phi(Z\widehat{\gamma})/\Phi(Z\widehat{\gamma})$ is formed for selected observations and merged onto the panel. Stage 2 is a linear panel regression on the selected sample with $\widehat{\lambda}$ added as a regressor. For each specification the Stage-2 model is estimated twice on a common Stage-1 inverse Mills ratio: "Entity" includes investor fixed effects only; "Entity+Time" adds month fixed effects. Province controls are absorbed by the investor fixed effect and are not separately included. Each variable spans two rows: the coefficient and, in parentheses below, the two-way (investor and month) cluster-robust t-statistic. This t is not the Newey–West / Fama–MacBeth statistic reported in the main tables; the estimator and the standard errors differ, so the notation deliberately differs. The inverse Mills ratio coefficient equals $\delta = \rho\sigma$ and is the control-function selection test: if insignificant, selection adds nothing beyond the fixed effects and the uncorrected fixed-effects estimator is consistent; if significant, its sign carries the direction of the selection bias. Risk preference, age, and the female indicator are time-invariant within investor and are absorbed by the entity fixed effect; they cannot be estimated in this design. Continuous regressors are winsorised at 1/99 and z-scored over the pooled sample (coefficients are per one pooled standard deviation), which differs from the within-month standardisation in the main Fama–MacBeth tables; levels are therefore not directly comparable across the two, though signs and significance are. All regressors enter at $t-1$ except the (absorbed) contemporaneous demographics. The +Biases panel survives in 35 of 50 months because its larger probit loses months to the pre-flight screens. The reported within R^2 is the relevant fit statistic; between and overall R^2 from the entity (and time) demeaning are out-of-transform and routinely negative, and are not reported. The Entity+Time within R^2 for the +Skills specification is omitted (marked n/a) as out-of-transform with the generated regressor. The Stage-2 clustered standard errors treat $\widehat{\lambda}$ as observed and do not account for its being a generated regressor; they understate uncertainty, most acutely on the inverse Mills ratio coefficient itself. Rigorous inference would require an investor-block bootstrap over both stages. The inverse Mills ratio t-statistics in particular should be read as approximate. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed).

Table A.6: Standalone Automated-Trading (Auto-Invest) Heckman Results

	BASE		SKILL		BIAS	
Variable	Selection ME	Outcome	Selection ME	Outcome	Selection ME	Outcome
Constant	-0.185***	1.619***	-0.177***	1.668***	-0.205***	1.190***
	(-8.840)	(17.826)	(-7.254)	(14.387)	(-7.689)	(66.054)
Log(TFA) (t-1)	0.122***	-0.305***	0.102***	-0.303***	0.057***	-0.173***
	(11.100)	(-9.439)	(9.780)	(-8.129)	(4.752)	(-11.584)
Risk Preference	0.015***	-0.025***	0.008**	-0.018***	0.012***	-0.005**
	(4.170)	(-4.825)	(2.473)	(-4.048)	(5.369)	(-2.026)
Age	0.030***	-0.024***	0.029***	-0.024***	0.030***	-0.007***
	(14.492)	(-5.706)	(13.596)	(-5.802)	(12.002)	(-3.132)
Female	0.000	0.032***	0.002	0.031***	-0.026***	0.039***
	(0.011)	(3.555)	(0.264)	(3.094)	(-9.515)	(19.958)
Combined RA Dummy (t-1)	0.079***	-0.093***	0.071***	-0.088***	0.058***	-0.018*
	(18.080)	(-18.587)	(12.955)	(-13.421)	(13.297)	(-1.882)
Baby Assets (t-1)	-0.038***	0.058***	-0.030***	0.054***	-0.012***	0.021***
	(-14.895)	(14.384)	(-31.046)	(12.818)	(-11.225)	(13.481)
Fund Experience (mos)			0.044***	-0.043***		
			(8.838)	(-5.444)		
Selection Skill			0.005**	-0.011***		
			(2.353)	(-2.833)		
Timing Rate			0.001	-0.001		
			(0.181)	(-0.115)		
Anchoring Bias					0.011	-0.006**
					(0.874)	(-2.552)
Total Turnover					-0.088***	-0.064***
					(-21.371)	(-15.355)
Attention Bias					-0.023***	0.000
					(-4.322)	(0.091)
Disposition Effect					-0.009***	0.006
					(-4.911)	(1.382)
Underdiversification (HHI)					-0.034***	0.088***
					(-18.115)	(24.700)
Extrapolation Bias					-0.008**	-0.003
					(-2.020)	(-0.820)
Gambling Bias					-0.005	-0.004
					(-1.153)	(-1.309)
Confirmation Bias					0.009***	0.018***
					(4.277)	(5.015)
Herding Bias					0.010***	-0.006***
					(3.287)	(-7.164)
Months Since Reg. (t-1)	0.030***		0.028***		0.026***	
	(16.686)		(14.240)		(13.069)	
Inverse Mills Ratio		-0.714***		-0.757***		-0.342***
		(-16.883)		(-13.292)		(-20.045)
Province Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2 / R^2	0.057	0.064	0.053	0.072	0.096	0.142
N	4,799,616	1,583,226	4,471,044	1,528,452	3,214,908	984,505

Notes. Heckman two-stage model with **automated-trading (Auto-Invest) participation** as the selection variable and the **auto-invest share** of fund purchases as the outcome. First stage: $\text{Auto}_{it} = 1$ if the investor used auto-invest purchase plans in month t. Outcome: auto_share_{it} = fraction of total fund purchases executed via auto-invest (conditional on Auto $= 1$). Three nested specifications mirror the main-paper Table 2: BASE (6 demographics), SKILL (adds Fund Experience, Selection Skill, Timing Rate), BIAS (adds 9 biases). Selection (first-stage) values are approximate marginal effects on the probability of auto-invest participation, $\text{ME}_k = \phi(\Phi^{-1}(\bar{P}))\hat{\beta}_k$, scaled per specification by $\bar{P} = 0.330$ (BASE), 0.342 (SKILL), 0.306 (BIAS) on the 4.9-million investor-month panel, with factors 0.362, 0.367, 0.351; outcome-stage coefficients are unscaled and t-statistics and significance are invariant to the scaling. Estimated monthly via Fama–MacBeth with Newey–West HAC t-statistics (max lag $= 3$) in parentheses. The BIAS sample drops to 35 months because bias construction requires longer fund-holding histories. Comparison with the RA Heckman (main Table 2) reveals the sign-flipping pattern that motivates Table 4 in the main paper: on **total turnover**, the coefficient is strongly negative for Auto (-0.064, $t = -15.4$) but strongly positive for RA ($+0.041$, $t = 6.2$ large). Underdiversification enters positively in both the Auto-Invest outcome equation (+0.088) and the RA outcome equation (+0.020 in Table 1 Panel 3). The Auto-Invest outcome coefficient is much larger in magnitude, and the within-investor-month stacking design in Table 4 confirms a relative tilt toward Auto-Invest. That is, concentrated investors delegate meaningfully to RA when they adopt, but they delegate relatively more to Auto-Invest in product comparison. These sign flips in the stacked-product analysis identify behavioral biases as driving the auto-vs.-RA choice in opposite directions. $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.10$.

Table A.7: Mutual Fund Performance Controls in the RA Heckman

DV (outcome): RA share = RA / (RA + MF) / *Selection:* $\text{RA}_{it} = 1$			
Fund-performance measure	Lag-1 month	Cumulative 3-month	Cumulative 6-month
Panel A: Outcome-equation coefficient on fund performance (RA share)			
Excess return vs CSI300	−0.002	−0.001	−0.001
	(−1.114)	(−0.358)	(−0.451)
β-adjusted excess vs CSI300	0.000	0.002*	0.002*
	(0.545)	(1.929)	(1.721)
Category-adjusted return	0.001	0.002**	0.003**
	(1.522)	(2.089)	(2.180)
Panel B: Selection-equation coefficient on fund performance (probit ME, Combined RA)			
Excess return vs CSI300	0.001	0.003	0.004
	(0.196)	(0.291)	(0.341)
β-adjusted excess vs CSI300	0.006	0.009	0.013
	(0.983)	(0.883)	(1.051)
Category-adjusted return	0.008	0.016**	0.022**
	(1.400)	(1.973)	(2.397)
Panel C: Heckman structural parameters (outcome equation)			
Inverse Mills Ratio λ (Exc CSI300)	0.090**	0.098**	0.088***
	(1.983)	(2.526)	(3.167)
Inverse Mills Ratio λ (β-adj)	0.006	0.026*	0.040**
	(0.424)	(1.646)	(2.185)
Inverse Mills Ratio λ (Cat-adj)	0.088**	0.098**	0.087***
	(2.101)	(2.552)	(3.069)
Demographic controls	Yes	Yes	Yes
Province dummies	Yes	Yes	Yes
T (months, Exc CSI300)	48	47	45
N outcome (Exc CSI300)	1,498,676	1,486,037	1,447,264
Pseudo R^2 / R^2 (outcome, Exc)	0.085	0.088	0.087

Notes. Heckman two-stage with Fama–MacBeth monthly estimation and Newey–West HAC t-statistics (max lag $= 3$) in parentheses. Each cell in Panels A and B is a separate Heckman specification: the three performance measures (rows) interact with three horizons (columns) to yield 9 specifications. Each specification includes the full demographic control set (Log(TFA), Risk Preference, Age, Female, Combined RA Dummy, Baby Assets; months-since-registration as the exclusion restriction in the selection equation) plus province dummies; only the performance coefficient is shown. *Performance measures*: (i) Excess Return vs CSI300 $= r_t^{\text{fund}} - r_t^{\text{CSI300}}$; (ii) β-adjusted excess $= r_t^{\text{fund}} - \hat{\beta}_i \cdot r_t^{\text{CSI300}}$; (iii) Category-adjusted $= r_t^{\text{fund}} - \bar{r}_t^{\text{peer}}$ (peer-group excess). All returns winsorised at 1%/99% and standardised within month. Sample size and T differ modestly across the three performance measures: β-adjusted specifications have shorter T (37, 36, 34 for L1/C3/C6 respectively) because β estimation requires a minimum history. Only the Excess-vs-CSI300 row's T and N are reported in the footer. *Headline.* Fund-performance coefficients in the RA share equation are economically tiny (magnitudes ≤ 0.003) across all 9 specifications. Only category-adjusted returns at the 3- and 6-month horizons reach borderline 5% significance, and even then the point estimates are below 0.003. The Heckman selection term $\hat{\lambda}$ is consistently positive and significant for Exc-CSI300 and Cat-adj specifications, indicating selection on unobservables that positively affect both RA participation and RA share. $^{***}\, p < 0.01$, $^{**}\, p < 0.05$, $^{*}\, p < 0.10$.

Table A.8: Robo-Advisor Performance Controls in the RA Heckman

DV (outcome): RA share = RA / (RA + MF) / *Selection:* $\text{RA}_{it} = 1$			
RA-performance measure	Lag-1 month	Cumulative 3-month	Cumulative 6-month
Panel A: Outcome-equation coefficient on RA performance (RA share)			
Excess return vs CSI300	0.003	0.001	−0.002
	(0.950)	(0.485)	(−0.835)
β-adjusted excess vs CSI300	0.002	−0.001	−0.005*
	(0.535)	(−0.275)	(−1.685)
Strategy-adjusted return	0.007***	0.008**	0.008**
	(2.946)	(2.541)	(2.387)
Demographic controls	Yes	Yes	Yes
Province dummies	Yes	Yes	Yes
T (months)	48	48	48
N outcome (Exc CSI300)	1,387,457	1,327,953	1,206,106
Pseudo R^2 / R^2 (outcome, Exc)	0.144	0.143	0.142

Notes. Same empirical design as Table A.7 but substituting **RA** return measures for fund return measures in the outcome equation. RA returns are observed only conditional on RA participation, so the RA-performance measure does not enter the selection equation (which is identical across the 9 specs because it has the same regressors as the baseline). Each cell is a separate Heckman specification; each contains the full demographic control set (Log(TFA), Risk Preference, Age, Female, Combined RA Dummy, Baby Assets; months-since-registration as exclusion restriction). Fama–MacBeth with Newey–West HAC t-statistics (max lag = 3) in parentheses. *Performance measures*: (i) Excess return vs CSI300; (ii) β-adjusted excess; (iii) Strategy-adjusted $= r_t^{\text{RA}} - \bar{r}_t^{\text{strategy peer}}$ where the peer group is all RA strategies of similar target-risk allocation. All returns winsorised at 1%/99% and standardised within month. *Headline.* As with fund performance (Table A.7), RA-return coefficients are small: Excess and β-adjusted returns have coefficients below $|0.005|$ across all horizons. Only **strategy-adjusted** RA returns have significantly positive coefficients (0.007–0.008), consistent with investors allocating more to RA *strategies* that beat their peer set — i.e., a within-RA selection effect — rather than to the RA channel as a whole beating the market. $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.10$.

Table A.9: Decision Authority and Product Choice: Pre-Use Horizon and Sample Robustness

Variable	Horizon	Product-step coefficients: DIRECT → AUTO (γ_T)	Product-step coefficients: AUTO → RA (γ_S)	Equality test: $H_0 : \gamma_T = \gamma_S$: $\gamma_T - \gamma_S$	$\chi^2(1)$	p-value
Panel A. Composite indices						
Skill composite, authority-bias composite, placebo-bias composite, and pre-use engagement.						
Mechanism variables						
Skill composite index	$L = 3$	0.001 (0.005)	0.008 (0.006)	-0.007	0.477	0.490
	$L = 9$	-0.004 (0.006)	0.016** (0.007)	-0.020	3.073	0.080
	$L = 12$	-0.008 (0.006)	0.016** (0.007)	-0.023	3.760	0.053
	$L = 6$, full panel	0.004 (0.005)	0.011* (0.006)	-0.006	0.409	0.522
Authority-bias composite index	$L = 3$	-0.005 (0.005)	-0.019*** (0.005)	0.015	2.338	0.126
	$L = 9$	-0.004 (0.006)	-0.023*** (0.006)	0.019	2.848	0.092
	$L = 12$	-0.002 (0.006)	-0.023*** (0.007)	0.020	2.993	0.084
	$L = 6$, full panel	0.000 (0.005)	-0.022*** (0.005)	0.022	5.085	0.024
Placebo-bias composite index	$L = 3$	-0.013** (0.006)	0.006 (0.006)	-0.019	3.107	0.078
	$L = 9$	-0.010 (0.007)	-0.001 (0.007)	-0.009	0.539	0.463
	$L = 12$	-0.014* (0.007)	0.004 (0.007)	-0.018	1.839	0.175
	$L = 6$, full panel	-0.009 (0.006)	0.003 (0.006)	-0.012	1.099	0.294
Pre-use login frequency	$L = 3$	-0.040*** (0.006)	0.011* (0.006)	-0.051	25.263	<0.001
	$L = 9$	-0.040*** (0.006)	0.015** (0.006)	-0.056	25.354	<0.001
	$L = 12$	-0.038*** (0.006)	0.012* (0.007)	-0.050	17.890	<0.001
	$L = 6$, full panel	-0.035*** (0.006)	0.013** (0.006)	-0.048	21.274	<0.001
Product-step intercepts						
Automation (c_T) / delegation (c_S)	$L = 3$	-0.596*** (0.018)	0.297*** (0.017)	—	—	—
	$L = 9$	-0.601*** (0.019)	0.325*** (0.018)	—	—	—
	$L = 12$	-0.622*** (0.019)	0.370*** (0.017)	—	—	—
	$L = 6$, full panel	-0.543*** (0.015)	0.287*** (0.016)	—	—	—
Baseline-control interactions						
Log total financial assets	$L = 3$	-0.023** (0.010)	0.084*** (0.008)	—	—	—
	$L = 9$	-0.035*** (0.011)	0.091*** (0.008)	—	—	—
	$L = 12$	-0.047*** (0.012)	0.095*** (0.009)	—	—	—
	$L = 6$, full panel	-0.057*** (0.006)	0.080*** (0.006)	—	—	—
Risk preference	$L = 3$	-0.013*** (0.004)	0.031*** (0.006)	—	—	—
	$L = 9$	-0.018*** (0.005)	0.034*** (0.006)	—	—	—
	$L = 12$	-0.024*** (0.005)	0.035*** (0.007)	—	—	—
	$L = 6$, full panel	-0.013*** (0.004)	0.030*** (0.006)	—	—	—
Age	$L = 3$	0.005 (0.004)	0.017*** (0.006)	—	—	—
	$L = 9$	-0.007 (0.005)	0.029*** (0.006)	—	—	—
	$L = 12$	-0.012** (0.005)	0.034*** (0.006)	—	—	—
	$L = 6$, full panel	0.001 (0.004)	0.022*** (0.005)	—	—	—
Female	$L = 3$	0.033*** (0.011)	-0.043*** (0.015)	—	—	—
	$L = 9$	0.038*** (0.013)	-0.043** (0.017)	—	—	—
	$L = 12$	0.047*** (0.013)	-0.054*** (0.018)	—	—	—
	$L = 6$, full panel	0.027** (0.011)	-0.036** (0.015)	—	—	—
Baby assets	$L = 3$	0.010* (0.005)	-0.005 (0.006)	—	—	—
	$L = 9$	0.014** (0.006)	-0.006 (0.006)	—	—	—
	$L = 12$	0.014** (0.006)	-0.005 (0.006)	—	—	—
	$L = 6$, full panel	0.020*** (0.005)	-0.006 (0.005)	—	—	—
Model fit						
Within R^2	$L = 3$: 0.2634; $L = 9$: 0.2719; $L = 12$: 0.2943; $L = 6$, full panel: 0.2470					
Overall R^2	$L = 3$: -0.6263; $L = 9$: -0.6032; $L = 12$: -0.5924; $L = 6$, full panel: -0.5673					
Within-X condition number	$L = 3$: 5.312; $L = 9$: 5.447; $L = 12$: 5.599; $L = 6$, full panel: 5.390					
Panel B. Disaggregated skill and authority biases						
Selection and timing ability; confirmation, turnover, speculative-turnover, herding, and disposition measures.						
Mechanism variables						
Selection ability	$L = 3$	—	—	—	—	—
	$L = 9$	-0.011* (0.006)	0.013** (0.006)	-0.024	5.309	0.021
	$L = 12$	-0.018*** (0.006)	0.012* (0.007)	-0.030	6.610	0.010
	$L = 6$, full panel	-0.004 (0.005)	0.012** (0.005)	-0.015	2.702	0.100
Timing ability	$L = 3$	—	—	—	—	—
	$L = 9$	0.002 (0.005)	0.010 (0.006)	-0.007	0.510	0.475
	$L = 12$	0.001 (0.006)	0.010 (0.006)	-0.010	0.751	0.386
	$L = 6$, full panel	0.008 (0.005)	0.003 (0.005)	0.005	0.265	0.607
Confirmation bias	$L = 3$	—	—	—	—	—
	$L = 9$	0.040*** (0.007)	0.007 (0.008)	0.033	5.395	0.020
	$L = 12$	0.045*** (0.008)	-0.004 (0.009)	0.049	10.449	0.001
	$L = 6$, full panel	0.029*** (0.007)	0.012* (0.007)	0.018	2.003	0.157
Total turnover	$L = 3$	—	—	—	—	—
	$L = 9$	0.038*** (0.008)	-0.036*** (0.008)	0.074	30.484	<0.001
	$L = 12$	0.047*** (0.008)	-0.041*** (0.009)	0.088	35.074	<0.001
	$L = 6$, full panel	0.032*** (0.007)	-0.031*** (0.006)	0.062	30.975	<0.001

Continued on next page

Table A.9 (continued)

Variable	Horizon	Product-step coefficients: DIRECT → AUTO (γ_T)	Product-step coefficients: AUTO → RA (γ_S)	Equality test: $\gamma_T - \gamma_S$	$\chi^2(1)$	p-value
Speculative turnover	$L = 3$	—	—	—	—	—
	$L = 9$	-0.036*** (0.006)	-0.008 (0.006)	-0.028	6.973	0.008
	$L = 12$	-0.037*** (0.006)	-0.005 (0.006)	-0.032	8.510	0.004
	$L = 6$, full panel	-0.026*** (0.005)	-0.012** (0.005)	-0.014	2.136	0.144
Herding bias	$L = 3$	—	—	—	—	—
	$L = 9$	-0.015** (0.007)	-0.002 (0.006)	-0.013	1.218	0.270
	$L = 12$	-0.017** (0.007)	0.001 (0.007)	-0.018	2.006	0.157
	$L = 6$, full panel	-0.012** (0.006)	-0.001 (0.006)	-0.011	1.106	0.293
Disposition effect	$L = 3$	—	—	—	—	—
	$L = 9$	0.006 (0.006)	-0.009 (0.006)	0.015	1.701	0.192
	$L = 12$	0.006 (0.006)	-0.011 (0.007)	0.017	1.918	0.166
	$L = 6$, full panel	0.004 (0.005)	-0.008 (0.006)	0.012	1.398	0.237
Product-step intercepts						
Automation (c_T) / delegation (c_S)	$L = 3$	—	—	—	—	—
	$L = 9$	-0.604*** (0.009)	0.333*** (0.011)	—	—	—
	$L = 12$	-0.627*** (0.007)	0.378*** (0.009)	—	—	—
	$L = 6$, full panel	-0.546*** (0.007)	0.295*** (0.012)	—	—	—
Baseline-control interactions						
Log total financial assets	$L = 3$	—	—	—	—	—
	$L = 9$	-0.035*** (0.011)	0.080*** (0.028)	—	—	—
	$L = 12$	-0.045*** (0.012)	0.083*** (0.028)	—	—	—
	$L = 6$, full panel	-0.058*** (0.010)	0.069*** (0.020)	—	—	—
Risk preference	$L = 3$	—	—	—	—	—
	$L = 9$	-0.021*** (0.006)	0.034*** (0.007)	—	—	—
	$L = 12$	-0.027*** (0.006)	0.034*** (0.007)	—	—	—
	$L = 6$, full panel	-0.016*** (0.006)	0.030*** (0.006)	—	—	—
Age	$L = 3$	—	—	—	—	—
	$L = 9$	-0.008 (0.006)	0.030*** (0.006)	—	—	—
	$L = 12$	-0.013** (0.006)	0.034*** (0.007)	—	—	—
	$L = 6$, full panel	0.001 (0.005)	0.022*** (0.006)	—	—	—
Female	$L = 3$	—	—	—	—	—
	$L = 9$	0.045** (0.020)	-0.053*** (0.020)	—	—	—
	$L = 12$	0.056*** (0.021)	-0.064*** (0.021)	—	—	—
	$L = 6$, full panel	0.034* (0.017)	-0.046** (0.018)	—	—	—
Baby assets	$L = 3$	—	—	—	—	—
	$L = 9$	0.014*** (0.005)	-0.004 (0.004)	—	—	—
	$L = 12$	0.014*** (0.005)	-0.004 (0.005)	—	—	—
	$L = 6$, full panel	0.019*** (0.004)	-0.004 (0.004)	—	—	—
Model fit						
Within R^2	$L = 3$: —; $L = 9$: 0.2749; $L = 12$: 0.2973; $L = 6$, full panel: 0.2495					
Overall R^2	$L = 3$: —; $L = 9$: -0.6025; $L = 12$: -0.5920; $L = 6$, full panel: -0.5660					
Within-X condition number	$L = 3$: —; $L = 9$: 5.591; $L = 12$: 5.739; $L = 6$, full panel: 5.533					
Panel C. Placebo biases						
Under-diversification, gambling, anchoring, attention, and extrapolation biases.						
Mechanism variables						
Under-diversification (HHI)	$L = 3$	—	—	0.084	67.654	<0.001
	$L = 9$	0.056*** (0.007)	-0.036*** (0.006)	0.092	64.257	<0.001
	$L = 12$	0.055*** (0.007)	-0.034*** (0.006)	0.089	53.878	<0.001
	$L = 6$, full panel	0.050*** (0.006)	-0.030*** (0.005)	0.080	58.697	<0.001
Gambling bias	$L = 3$	—	—	0.003	0.097	0.756
	$L = 9$	0.002 (0.007)	-0.015** (0.007)	0.017	1.709	0.191
	$L = 12$	-0.001 (0.008)	-0.007 (0.007)	0.005	0.153	0.695
	$L = 6$, full panel	0.000 (0.006)	-0.012** (0.006)	0.012	1.211	0.271
Anchoring bias	$L = 3$	—	—	-0.087	67.767	<0.001
	$L = 9$	-0.044*** (0.007)	0.053*** (0.006)	-0.097	61.550	<0.001
	$L = 12$	-0.049*** (0.007)	0.063*** (0.007)	-0.111	66.223	<0.001
	$L = 6$, full panel	-0.036*** (0.006)	0.045*** (0.006)	-0.081	57.155	<0.001
Attention bias	$L = 3$	—	—	-0.064	32.332	<0.001
	$L = 9$	-0.042*** (0.007)	0.023*** (0.007)	-0.065	25.254	<0.001
	$L = 12$	-0.041*** (0.008)	0.021*** (0.007)	-0.062	19.683	<0.001
	$L = 6$, full panel	-0.041*** (0.006)	0.025*** (0.006)	-0.066	31.635	<0.001
Extrapolation bias	$L = 3$	—	—	0.023	4.242	0.039
	$L = 9$	0.006 (0.007)	-0.019*** (0.007)	0.025	4.070	0.044
	$L = 12$	0.004 (0.008)	-0.020*** (0.007)	0.024	2.958	0.086
	$L = 6$, full panel	0.010 (0.006)	-0.018*** (0.006)	0.028	6.045	0.014
Product-step intercepts						
Automation (c_T) / delegation (c_S)	$L = 3$	—	—	—	—	—
	$L = 9$	-0.618*** (0.019)	0.340*** (0.020)	—	—	—
	$L = 12$	-0.639*** (0.019)	0.387*** (0.019)	—	—	—

Continued on next page

Table A.9 (continued)

Variable	Horizon	Product-step coefficients: DIRECT → AUTO (γ_T)	Product-step coefficients: AUTO → RA (γ_S)	Equality test: $H_0 : \gamma_T = \gamma_S$: $\gamma_T - \gamma_S$	$\chi^2(1)$	p-value
	$L = 6$, full panel	-0.556*** (0.015)	0.299*** (0.019)	—	—	—
Baseline-control interactions						
Log total financial assets	$L = 3$	—	—	—	—	—
	$L = 9$	-0.011 (n.a.)	0.069*** (0.013)	—	—	—
	$L = 12$	-0.023 (n.a.)	0.070*** (0.013)	—	—	—
	$L = 6$, full panel	-0.039*** (0.001)	0.064*** (0.010)	—	—	—
Risk preference	$L = 3$	—	—	—	—	—
	$L = 9$	-0.018*** (0.005)	0.034*** (0.006)	—	—	—
	$L = 12$	-0.023*** (0.005)	0.034*** (0.007)	—	—	—
	$L = 6$, full panel	-0.013*** (0.004)	0.030*** (0.005)	—	—	—
Age	$L = 3$	—	—	—	—	—
	$L = 9$	-0.005 (0.005)	0.028*** (0.006)	—	—	—
	$L = 12$	-0.009* (0.005)	0.032*** (0.006)	—	—	—
	$L = 6$, full panel	0.003 (0.004)	0.021*** (0.005)	—	—	—
Female	$L = 3$	—	—	—	—	—
	$L = 9$	0.057*** (0.009)	-0.057*** (0.013)	—	—	—
	$L = 12$	0.066*** (0.009)	-0.069*** (0.013)	—	—	—
	$L = 6$, full panel	0.043*** (0.009)	-0.047*** (0.012)	—	—	—
Baby assets	$L = 3$	—	—	—	—	—
	$L = 9$	0.011** (0.005)	-0.003 (0.005)	—	—	—
	$L = 12$	0.011** (0.005)	-0.003 (0.005)	—	—	—
	$L = 6$, full panel	0.017*** (0.005)	-0.005 (0.005)	—	—	—
Model fit						
Within R^2	$L = 3$: —; $L = 9$: 0.2758; $L = 12$: 0.2991; $L = 6$, full panel: 0.2505					
Overall R^2	$L = 3$: —; $L = 9$: -0.6035; $L = 12$: -0.5925; $L = 6$, full panel: -0.5674					
Within-X condition number	$L = 3$: —; $L = 9$: 5.558; $L = 12$: 5.724; $L = 6$, full panel: 5.477					
Panel D. Sample construction						
Stacked observations	$L = 3$: 766,254; $L = 9$: 593,190; $L = 12$: 496,575; $L = 6$, full panel: 728,499					
Investor-month choice sets	$L = 3$: 255,418; $L = 9$: 197,730; $L = 12$: 165,525; $L = 6$, full panel: 242,833					
Unique investors	$L = 3$: 12,839; $L = 9$: 10,884; $L = 12$: 9,707; $L = 6$, full panel: 11,790					
Calendar months	$L = 3$: 35; $L = 9$: 33; $L = 12$: 31; $L = 6$, full panel: 34					
Choice-set fixed effects	Investor-month (investor × calendar month); no separate month effects					
Clustering	Two-way, by investor and by calendar month					
Pre-use window	$[t_{\text{first}} - L,\ t_{\text{first}} - 1]$					
Eligibility threshold	At least $\max(2, L/2)$ non-missing pre-use months per type variable					
Type-variable scaling	Per-investor pre-use averages, then cross-sectional z-scores					

Notes. This table reports the sensitivity of the decision-authority estimates of Table 4 to the horizon L over which investor type is measured and to the sample restriction. Within each variable block, the first three rows vary the pre-use horizon to $L = 3$, 9, and 12 months while holding the identifying-only restriction of the headline model fixed; the fourth row holds the horizon at the headline value $L = 6$ and instead retains all eligible investor-months. The headline model itself ($L = 6$, identifying-only) is reported in Table 4 and is not repeated here, so the horizon ladder reads $L = 3$, then the headline $L = 6$ in Table 4, then $L = 9$ and $L = 12$. This table also reports the baseline-control interactions that are suppressed in Table 4.

Each investor-month is stacked into three rows—DIRECT (self-directed mutual fund investing), AUTO (automatic investment plans, where the investor selects funds and the platform automates purchase timing), and RA (the robo-advisor, where the algorithm selects funds and manages the portfolio)—with the binary use indicator as the dependent variable. Products may be held simultaneously, so this is a stacked multivariate linear probability model rather than a mutually exclusive multinomial choice model. The cumulative-authority indicator c_T equals one for AUTO and RA, and c_S equals one for RA only. For a characteristic X, γ_T is the coefficient on $X \cdot c_T$ (the *automation step*, DIRECT → AUTO, where execution is delegated but fund choice is retained) and γ_S is the coefficient on $X \cdot c_S$ (the *full-delegation step*, AUTO → RA, where fund choice is surrendered).

Cell format. The γ_T and γ_S columns report the coefficient with significance stars, followed by the two-way cluster-robust standard error in parentheses. The t-statistic is not shown separately because it is exactly the coefficient divided by that standard error, and the stars are assigned from it. Equality columns report $\Delta = \gamma_T - \gamma_S$, the Wald statistic $\chi^2(1)$ for $H_0 : \gamma_T = \gamma_S$, and its p-value; equality tests are not defined for the product-step intercepts or the baseline controls, which are conditioning terms rather than objects of interest, and those cells are marked —. Cells marked — under $L = 3$ in Panels B and C are coefficient blocks that were not supplied with the estimation output; where the corresponding equality test was supplied it is reported. In Panel C the estimation returns the log-assets × automation coefficient but no clustered standard error under $L = 9$ and $L = 12$; those coefficients are reported without significance stars and their standard errors are marked (n.a.).

Type measures are per-investor averages over the L months immediately preceding the investor's first AUTO or RA use, requiring at least $\max(2, L/2)$ non-missing pre-use months for every type variable, and are z-scored cross-sectionally; composite indices are formed from the z-scored components. Estimation is by `linearmodels.PanelOLS` with entity fixed effects on the investor-month. Standard errors are two-way clustered by investor and by calendar month. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed). Within R^2 is the relevant goodness-of-fit measure under the fixed-effects transformation; negative overall R^2 is a known artefact of fixed-effects models and does not indicate negative in-sample fit.

The sample shrinks as the horizon lengthens, because a longer pre-use window requires a longer observed history before first use and therefore removes both early-adopting investors and early calendar months: stacked observations fall from 766,254 at $L = 3$ to 593,190 at $L = 9$ and 496,575 at $L = 12$, and calendar months from 35 to 33 and 31. Comparisons across horizons are therefore across partly different investors, and the $L = 12$ column in particular rests on 9,707 investors against 12,839 at $L = 3$. These estimates are descriptive mechanism evidence and should be read as associations consistent with the decision-authority mechanism, not as causal effects.

Table A.10: Pre-RA Experience and Robo-Advisor Allocation over Tenure: Robustness to Adopter-Sample Restrictions

Variable	Consecutive RA-months only	No-gap stayers	Non-zero tenure
Pre-RA experience and tenure interactions			
Pre-RA experience	-0.181***	-0.173***	-0.181***
	(-3.007)	(-2.820)	(-3.016)
Pre-RA experience × tenure	-0.036**	-0.040*	-0.036**
	(-1.991)	(-1.877)	(-2.001)
Pre-RA experience × tenure^2	-0.013	-0.018	-0.013
	(-1.218)	(-1.209)	(-1.218)
Tenure (main effects)			
Tenure (within-month demeaned)	-0.042**	-0.045**	-0.042**
	(-2.150)	(-2.042)	(-2.153)
Tenure^2	-0.013	-0.018	-0.012
	(-1.241)	(-1.233)	(-1.241)
Total observations	1,049,512	1,036,734	1,054,184
R^2	0.157	0.159	0.158

Notes. This table reproduces the quadratic OLS specification of Table 8 (column 2) under three alternative adopter-sample restrictions. The dependent variable, the construction of *Pre-RA experience* and *Tenure*, the baseline controls (which are included in every regression but suppressed here), the Fama-MacBeth estimator, and the Newey-West HAC standard errors are as in Table 8.

The three samples differ as follows. *Consecutive RA-months only* (column 1) retains an investor-month only if the prior month was also robo-advisor-active, isolating periods of continuous tenure. *No-gap stayers* (column 2) retains only those adopters whose robo-advisor-active history contains no gap between the first adoption month and the end of the sample window; non-stayers are excluded entirely. *Non-zero tenure* (column 3) drops the first robo-advisor-active month per investor, so the regression identifies the relation between the focal regressor and within-tenure variation in allocation; the very first month, which mechanically has zero tenure, is removed.

Newey-West HAC t-statistics in parentheses below coefficient estimates. Significance levels (against standard-normal critical values): *** $|t| \geq 2.576$ ($p < 0.01$); ** $|t| \geq 1.960$ ($p < 0.05$); * $|t| \geq 1.645$ ($p < 0.10$); two-sided.

Table A.11: Additional Bias Trajectories over Robo-Advisor Tenure

Regressor	Attention (att)	Disposition (dep)	Gambling (gam)	Herding (her)
Panel A: Rolling 6-month trailing mean of bias (headline)				
Within-adopter sample				
τ	$+0.0003$	-0.0010^{***}	-0.0002	$+0.0003^{*}$
	$(+0.81)$	(-3.81)	(-1.18)	$(+1.73)$
$\tau\times$ High Under-div.	$+0.0016^{***}$	$+0.0009$	$+0.0002$	$+0.0017^{***}$
	$(+3.31)$	$(+1.46)$	$(+0.80)$	$(+4.38)$
$\tau\times$ High Turnover	$+0.0015^{***}$	$+0.0008$	$+0.0008^{***}$	$+0.0015^{***}$
	$(+3.29)$	$(+1.64)$	$(+3.25)$	$(+3.88)$
$\tau\times$ AOT	$+0.0013^{***}$	$+0.0006$	$+0.0004^{*}$	$+0.0010^{***}$
	$(+3.46)$	$(+1.60)$	$(+1.73)$	$(+3.49)$
Difference-in-differences with matched never-adopters				
Adopter $\times\ \tau$	$+0.0001$	-0.0006^{***}	-0.0006^{***}	$+0.0004^{***}$
	$(+0.55)$	(-3.21)	(-4.30)	$(+2.76)$
Adopter $\times\ \tau\times$ Engagement	$+0.0003^{***}$	$+0.0003^{***}$	$+0.0002^{***}$	$+0.0001^{**}$
	$(+3.06)$	$(+3.25)$	$(+2.81)$	$(+2.08)$
Adopter $\times\ \tau\times$ High Under-div.	$+0.0017^{***}$	$+0.0008$	$+0.0003$	$+0.0016^{***}$
	$(+3.49)$	$(+1.33)$	$(+0.90)$	$(+4.37)$
Adopter $\times\ \tau\times$ High Turnover	$+0.0015^{***}$	$+0.0008$	$+0.0008^{***}$	$+0.0015^{***}$
	$(+3.25)$	$(+1.64)$	$(+3.28)$	$(+3.89)$
Adopter $\times\ \tau\times$ AOT	$+0.0013^{***}$	$+0.0006$	$+0.0004^{*}$	$+0.0010^{***}$
	$(+3.52)$	$(+1.53)$	$(+1.81)$	$(+3.49)$
Observations (adopter-only)	1,321,291	1,300,247	1,321,291	1,321,291
Observations (DiD sample)	1,847,939	1,801,359	1,847,939	1,847,939
Panel B: Raw monthly bias (robustness)				
Within-adopter sample				
τ	-0.0000	-0.0010^{***}	-0.0002	-0.0001
	(-0.06)	(-3.28)	(-0.70)	(-0.26)
$\tau\times$ High Under-div.	$+0.0014^{***}$	$+0.0004$	$+0.0003$	$+0.0008^{**}$
	$(+2.58)$	$(+0.66)$	$(+1.12)$	$(+2.55)$
$\tau\times$ High Turnover	$+0.0008$	$+0.0002$	$+0.0009^{***}$	$+0.0007^{**}$
	$(+1.59)$	$(+0.56)$	$(+3.01)$	$(+2.16)$
$\tau\times$ AOT	$+0.0009^{**}$	$+0.0001$	$+0.0005^{**}$	$+0.0004^{*}$
	$(+2.12)$	$(+0.36)$	$(+2.00)$	$(+1.80)$
Difference-in-differences with matched never-adopters				
Adopter $\times\ \tau$	$+0.0005$	-0.0006^{***}	-0.0008^{***}	-0.0001
	$(+1.36)$	(-3.23)	(-4.99)	(-0.69)
Adopter $\times\ \tau\times$ Engagement	$+0.0001$	$+0.0002^{*}$	$+0.0001^{**}$	$+0.0001^{*}$
	$(+1.26)$	$(+1.74)$	$(+2.15)$	$(+1.86)$
Adopter $\times\ \tau\times$ High Under-div.	$+0.0014^{***}$	$+0.0003$	$+0.0004$	$+0.0008^{***}$
	$(+2.58)$	$(+0.57)$	$(+1.30)$	$(+2.58)$
Adopter $\times\ \tau\times$ High Turnover	$+0.0008$	$+0.0002$	$+0.0009^{***}$	$+0.0007^{**}$
	$(+1.57)$	$(+0.56)$	$(+3.01)$	$(+2.17)$
Adopter $\times\ \tau\times$ AOT	$+0.0009^{**}$	$+0.0001$	$+0.0005^{**}$	$+0.0005^{*}$
	$(+2.10)$	$(+0.31)$	$(+2.08)$	$(+1.82)$
Observations (adopter-only)	1,274,140	1,309,740	1,274,140	1,274,140
Observations (DiD sample)	1,781,292	1,815,382	1,781,292	1,781,292

Notes. The table extends the bias-trajectory analysis to four additional behavioural bias measures: attention (att), disposition effect (dep), gambling preference (gam), and herding (her). The specification, dependent-variable construction, sample, control set, fixed-effects structure, and standard-error treatment are identical to those in the main bias-trajectory table; see the notes there for full detail. Panel A uses the rolling six-month trailing mean of the bias measure as the dependent variable; Panel B replaces this with the raw monthly bias as a robustness check on the smoothing window.

These four biases are not treated as headline focals or as clean placebos because they admit alternative interpretations or fail to deliver a clean signal on the engagement dose-response test.

Attention (att) and gambling (gam) are effectively null on the core falsification tests. The Adopter $\times\ \tau$ coefficient is statistically indistinguishable from zero for attention in both panels and is small and unstable in sign for gambling across the two panels. The Adopter $\times\ \tau\times$ Engagement triple is

marginally significant for attention in Panel A ($t = +3.06$) but shrinks to insignificance in the raw monthly Panel B ($t = +1.26$), which is consistent with a smoothing-window artifact rather than a genuine engagement effect. The pattern supports treating these biases as additional placebos rather than as channels of selective learning.

Disposition effect (dep) shows a small negative Adopter $\times \tau$ in both panels that is robust to the smoothing window ($t = -3.21$ in Panel A and $t = -3.23$ in Panel B). The economic interpretation is mechanical rather than psychological. The disposition effect is constructed from investor-initiated realisations of gains and losses on the mutual-fund portfolio. Robo-advisor adoption moves a portion of investor wealth into a delegated portfolio that the investor does not actively trade, so the share of investor activity available to generate measurable realisations falls mechanically with tenure. The negative coefficient on Adopter $\times \tau$ therefore reflects a denominator effect on the disposition statistic rather than a direct change in psychological loss aversion, and we treat it as non-diagnostic for the selective-learning channel.

Herding (her) rises positively with tenure across both panels (Adopter $\times \tau$ has $t = +2.76$ in Panel A; the raw counterpart in Panel B is insignificant at $t = -0.69$ but turns positive once interactions with high-decile focal-bias indicators are added). The natural interpretation is that delegating wealth to a robo-advisor exposes the investor to the strategy's portfolio adjustments and trades, which the investor then partially mirrors in the self-directed mutual-fund portfolio. The directional opposite of the prediction for the focal biases is therefore not a falsification of the selective-learning channel, but a co-occurring tracking channel that operates through observation of the robo-advisor's holdings rather than through investor re-evaluation of own beliefs.

Variable definitions match the main bias-trajectory table. τ is months since first robo-advisor adoption, capped at 24. High Under-div. and High Turnover are top-decile indicators on the six-month pre-adoption mean of the corresponding focal bias. AOT (active opinion-taker) equals one if either focal high-decile indicator is one. Engagement is the standardised lagged 12-month login count. All specifications include lagged log total financial assets, lagged cash-management asset holdings, and lagged standardised login count as level controls; coefficients on these controls are omitted from the table.

Significance levels are based on standard-normal critical values applied to cluster-robust t-statistics: *** $|t| \geq 2.576$ ($p < 0.01$); ** $|t| \geq 1.960$ ($p < 0.05$); * $|t| \geq 1.645$ ($p < 0.10$); two-sided. Cluster-robust t-statistics are reported in parentheses below coefficient estimates.

Table A.12: First Robo-Advisor Adoption: Heckman Two-Step Fama–MacBeth, Fund-Return Predictors

	(1) Lag 1 month		(2) Cumulative 3 months		(3) Cumulative 6 months	
Variable	Stage 1 ME	Stage 2	Stage 1 ME	Stage 2	Stage 1 ME	Stage 2
Panel A: Pre-window construction (adopter type)						
Fund excess CSI300	0.130	0.132	-0.027	-0.246	-0.004	-0.059*
	(1.399)	(0.658)	(-0.673)	(-1.412)	(-0.283)	(-1.741)
Fund CAPM alpha (CSI300)	-0.127	-0.132	0.030	0.229	0.011	0.035
	(-1.328)	(-0.687)	(0.767)	(1.268)	(0.627)	(1.409)
Fund category-adjusted	-0.002	-0.000	-0.000	0.011***	-0.003	0.009
	(-0.926)	(-0.098)	(-0.037)	(2.712)	(-0.665)	(1.260)
Months since reg., $t-1$	-0.003		-0.003		-0.003	
	(-1.020)		(-1.145)		(-1.131)	
Inverse Mills ratio		0.061*		0.014		-0.012
		(1.766)		(0.261)		(-0.341)
Pseudo R^2 / R^2	0.210	—	0.229	—	0.245	—
T (months)	15	15	17	17	19	19
Avg. N per month (S1 / S2)	8,079 / 796		8,761 / 769		9,212 / 748	
Panel B: Lag-1 hazard-style design (adoption hazard)						
Fund excess CSI300	0.024	-0.007	0.002	0.001	0.002	-0.014
	(1.183)	(-0.345)	(1.455)	(0.071)	(1.524)	(-1.543)
Fund CAPM alpha (CSI300)	-0.023	0.012	-0.001*	-0.014	-0.001	0.003
	(-1.119)	(0.552)	(-1.842)	(-1.030)	(-1.066)	(0.500)
Fund category-adjusted	-0.000	0.008	0.001*	0.011**	0.001**	0.001
	(-0.581)	(0.772)	(1.725)	(2.180)	(2.446)	(0.193)
Months since reg., $t-1$	0.002***		0.002***		0.002***	
	(3.550)		(4.011)		(3.663)	
Inverse Mills ratio		-0.076		-0.026		-0.134**
		(-0.760)		(-0.273)		(-2.439)
Pseudo R^2 / R^2	0.060	—	0.060	—	0.059	—
T (months)	25	25	25	25	27	27
Avg. N per month (S1 / S2)	34,402 / 840		34,348 / 829		34,140 / 791	
Baseline + Skills + Biases	Included	Included	Included	Included	Included	Included
Province Dummies	Yes	Yes	Yes	Yes	Yes	Yes

Notes. Estimates are from unweighted Heckman two-step (Tobit-2) selection models estimated monthly and aggregated by Fama–MacBeth. The selection (Stage 1) dependent variable is $D_{i,t} = 1$ if month t is the investor's first robo-advisor month; the outcome (Stage 2) dependent variable is the at-adoption share $rsh_{i,t} = RA_{i,t}/(RA_{i,t}+MF_{i,t})$ on the $D = 1$ cross-section. Three fund-return regressors are added on top of the Baseline + Skills + Biases specification (full coefficients reported in the main first-adoption table): fund excess return over the CSI 300, fund CAPM alpha relative to the CSI 300, and fund category-adjusted return. Each column corresponds to one horizon. Returns enter both stages because they are mutual-fund returns, defined for every at-risk investor. Panel A uses the pre-window construction (per-investor means over the entire pre-adoption history, z-scored within month); Panel B uses the lag-1 hazard-style construction (regressor read at calendar month $t-1$, winsorised and z-scored within month). The two designs are complementary: the pre-window identifies adopter type, the lag-1 identifies the timing of the adoption hazard. The exclusion restriction is months-since-registration ($t-1$), in Stage 1 only. Stage-1 (selection) values are approximate marginal effects on the first-adoption probability, $\mathrm{ME}_k = \phi(\Phi^{-1}(\bar{P}))\hat{\beta}_k$; because the first-adoption rate differs by horizon, each column is scaled by its own rate – Panel A factors 0.174, 0.159, 0.150 and Panel B factors 0.057, 0.057, 0.055 (lag-1/cum-3/cum-6), on the 4.9-million investor-month at-risk panel. Stage-2 outcome coefficients are unscaled; t-statistics and significance are invariant to the scaling. Newey–West HAC t-statistics (lag = min$(3, T-1)$) in parentheses below estimates. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed). Province dummies included, not reported.

Table A.13: First Auto-Invest Adoption: Heckman Two-Step Fama–MacBeth, Fund-Return Predictors

	(1) Lag 1 month		(2) Cumulative 3 months		(3) Cumulative 6 months	
Variable	Stage 1 ME	Stage 2	Stage 1 ME	Stage 2	Stage 1 ME	Stage 2
Panel A: Pre-window construction (adopter type)						
Fund excess CSI300	0.025	-0.262	0.005	0.024	-0.009	-0.113**
	(0.590)	(-1.290)	(0.153)	(0.141)	(-0.883)	(-2.045)
Fund CAPM alpha (CSI300)	-0.019	0.264	0.008	-0.038	0.016*	0.077
	(-0.425)	(1.350)	(0.217)	(-0.206)	(1.712)	(1.320)
Fund category-adjusted	-0.003	0.015	-0.004	0.005	0.002	0.018
	(-0.507)	(0.683)	(-1.113)	(0.305)	(0.468)	(1.087)
Months since reg., $t-1$	-0.003		-0.003		-0.004	
	(-1.060)		(-1.189)		(-1.455)	
Inverse Mills ratio		0.092		0.008		0.005
		(0.902)		(0.064)		(0.035)
Pseudo R^2 / R^2	0.247	—	0.237	—	0.251	—
T (months)	20	20	22	22	23	23
Avg. N per month (S1 / S2)	1,931 / 198		1,720 / 167		1,797 / 169	
Panel B: Lag-1 hazard-style design (adoption hazard)						
Fund excess CSI300	0.001	0.026	-0.003	0.061	-0.001	0.320**
	(0.930)	(0.082)	(-0.980)	(0.501)	(-0.454)	(2.044)
Fund CAPM alpha (CSI300)	-0.000	0.054	0.003	0.008	0.001	-0.364**
	(-0.297)	(0.183)	(0.977)	(0.065)	(0.310)	(-2.380)
Fund category-adjusted	-0.001*	0.129**	0.001**	0.010	0.001***	0.052
	(-1.767)	(1.985)	(2.109)	(0.505)	(3.902)	(1.292)
Months since reg., $t-1$	0.001***		0.001***		0.001**	
	(2.863)		(3.667)		(2.284)	
Inverse Mills ratio		-0.517		-0.912		-0.806
		(-0.538)		(-1.245)		(-0.428)
Pseudo R^2 / R^2	0.039	—	0.035	—	0.032	—
N	310,125	6,660	392,331	6,840	331,328	5,776
Baseline + Skills + Biases + Holdings + Engagement	Included		Included		Included	
Province Dummies	Yes	Yes	Yes	Yes	Yes	Yes

Notes. Estimates are from unweighted Heckman two-step (Tobit-2) selection models estimated monthly and aggregated by Fama–MacBeth. The selection (Stage 1) dependent variable is $D_{i,t} = 1$ if month t is the investor's first auto-invest month; the outcome (Stage 2) dependent variable is the at-adoption AUTO purchase share $ash_{i,t} = auto_buy_{i,t}/total_buy_{i,t}$ on the $D = 1$ cross-section. Three fund-return regressors are added on top of the full pre-return specification (Baseline + Skills + Biases + Holdings + Engagement; coefficients on the controls are stable across the return horizons and omitted): fund excess return over the CSI 300, fund CAPM alpha relative to the CSI 300, and fund category-adjusted return. Each column corresponds to one horizon. Returns enter both stages because they are mutual-fund returns, defined for every at-risk investor. Panel A uses the pre-window construction (per-investor means over the entire pre-adoption history, z-scored within month); Panel B uses the lag-1 hazard-style construction (regressor read at calendar month $t-1$, winsorised and z-scored within month). The two designs are complementary: the pre-window identifies adopter type, the lag-1 identifies the timing of the adoption hazard. The exclusion restriction is months-since-registration ($t-1$), in Stage 1 only. Stage-1 (selection) values are approximate marginal effects on the first-adoption probability, $\mathrm{ME}_k = \phi(\Phi^{-1}(\bar{P}))\hat{\beta}_k$; each column is scaled by its own rate – Panel A factors 0.179, 0.172, 0.168 and Panel B factors 0.044, 0.043, 0.043 (lag-1/cum-3/cum-6), on the 4.9-million investor-month at-risk panel. Stage-2 outcome coefficients are unscaled; t-statistics and significance are invariant to the scaling. Newey–West HAC t-statistics (lag = $\min(3, T-1)$) in parentheses below estimates. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed). Province dummies included, not reported.

Table A.14: Return Predictors of Robo-Advisor Allocation across Adoption-Lifecycle Margins: Full Specifications

	Horizon		
Variable	Lag 1 month	Cumulative 3 months	Cumulative 6 months
Panel A: First adoption (Stage 2 outcome)			
A.1: Positive/negative decomposition, raw returns			
Fund return^{+}	0.003**	0.003**	0.009***
	(2.066)	(2.496)	(4.196)
Fund return^{-}	0.005*	0.006**	0.012***
	(1.956)	(2.130)	(3.474)
Inverse Mills ratio	0.378	0.172	-0.104
	(1.242)	(0.652)	(-0.635)
A.2: Positive/negative decomposition, excess over CSI 300			
Fund excess^{+}	0.006***	0.006***	0.009***
	(3.625)	(3.355)	(3.551)
Fund excess^{-}	0.010***	0.013***	0.015***
	(4.335)	(4.384)	(4.983)
Inverse Mills ratio	0.410	0.157	-0.106
	(1.266)	(0.644)	(-0.641)
A.3: Above-median indicator, raw returns			
Fund: above median	-0.001	-0.003	-0.005
	(-0.178)	(-0.880)	(-0.844)
Inverse Mills ratio	0.368	0.170	-0.064
	(1.218)	(0.658)	(-0.413)
A.4: Above-median indicator, excess over CSI 300			
Fund excess: above median	-0.001	-0.004	-0.005
	(-0.177)	(-0.987)	(-0.949)
Inverse Mills ratio	0.369	0.166	-0.073
	(1.219)	(0.646)	(-0.459)
Observations (uncensored, A.1)	51,370	51,921	53,764
T (months, A.1)	34	34	36
Panel B: Exit (Stage 1 selection, probit-only; marginal effects)			
B.1: Positive/negative decomposition, raw returns			
Fund return^{+}	0.002***	0.004***	0.004***
	(2.783)	(4.703)	(8.644)
Fund return^{-}	-0.003***	-0.003***	-0.002**
	(-3.017)	(-2.912)	(-2.306)
RA return^{+}	0.001	-0.001	0.000
	(0.777)	(-0.562)	(0.262)
RA return^{-}	0.003***	0.001	0.001
	(2.928)	(0.688)	(0.986)
B.2: Positive/negative decomposition, excess over CSI 300			
Fund excess^{+}	0.005***	0.006***	0.007***
	(3.300)	(4.746)	(5.755)
Fund excess^{-}	0.001	0.002*	0.002*
	(1.284)	(1.746)	(1.889)
RA excess^{+}	-0.001	-0.000	0.001
	(-0.455)	(-0.070)	(1.066)
RA excess^{-}	0.000	0.001	0.002**
	(0.065)	(0.383)	(2.150)
B.3: Above-median indicator, raw returns			
Fund: above median	0.006***	0.009***	0.009***
	(3.241)	(4.983)	(5.276)
RA: above median	-0.001	0.000	0.001
	(-0.578)	(0.162)	(0.490)
B.4: Above-median indicator, excess over CSI 300			
Fund excess: above median	0.006***	0.009***	0.009***
	(3.217)	(5.028)	(5.710)
RA excess: above median	-0.001	0.001	0.002
	(-0.582)	(0.242)	(1.067)
Observations (B.1)	1,263,735	1,177,032	1,118,619
T (months, B.1)	42	40	41

Continued on next page

Table A.14 (continued)

	Horizon		
Variable	Lag 1 month	Cumulative 3 months	Cumulative 6 months
Panel C: Continuation (Stage 2 outcome)			
C.1: Positive/negative decomposition, raw returns			
Fund return$^+$	0.012***	0.013***	0.015***
	(12.824)	(12.251)	(9.188)
Fund return$^-$	0.011***	0.011***	0.013***
	(8.246)	(6.517)	(6.034)
RA return$^+$	0.007**	0.009***	0.009***
	(2.370)	(2.869)	(3.119)
RA return$^-$	0.002	0.005	0.006**
	(0.575)	(1.629)	(2.338)
Inverse Mills ratio	0.333**	0.396***	0.398***
	(2.292)	(2.669)	(2.691)
C.2: Positive/negative decomposition, excess over CSI 300			
Fund excess$^+$	0.012***	0.014***	0.017***
	(12.874)	(10.159)	(6.312)
Fund excess$^-$	0.015***	0.016***	0.016***
	(9.175)	(7.852)	(9.101)
RA excess$^+$	0.000	-0.004	-0.005
	(0.108)	(-1.620)	(-1.499)
RA excess$^-$	-0.003	-0.004*	-0.002
	(-1.011)	(-1.888)	(-0.903)
Inverse Mills ratio	0.309**	0.394**	0.399***
	(2.107)	(2.569)	(2.746)
C.3: Above-median indicator, raw returns			
Fund: above median	0.001	0.003	0.003
	(0.313)	(1.201)	(0.919)
RA: above median	0.001	0.001	-0.003
	(0.296)	(0.335)	(-0.512)
Inverse Mills ratio	0.301**	0.382**	0.406***
	(2.020)	(2.486)	(2.692)
C.4: Above-median indicator, excess over CSI 300			
Fund excess: above median	0.001	0.002	0.002
	(0.361)	(1.061)	(0.699)
RA excess: above median	0.001	0.001	-0.001
	(0.319)	(0.147)	(-0.238)
Inverse Mills ratio	0.302**	0.384**	0.408***
	(2.027)	(2.516)	(2.748)
N	1,328,226	1,271,406	1,147,071
Baseline controls	Yes	Yes	Yes
Province dummies	Yes	Yes	Yes

Notes. Estimates are from unweighted Heckman two-step (Tobit-2) models estimated monthly and aggregated by Fama–MacBeth. The outcome dependent variable is $ra_share_{i,t} = RA_{i,t}/(RA_{i,t} + MF_{i,t})$, observed only when the investor uses the robo-advisor. The selection equation defines a different at-risk margin in each panel: Panel A restricts to investor-months up to and including the first robo-advisor month, with $D = 1$ in that adoption month; Panel B restricts to investor-months held in $t-1$ and sets $D = 1$ if the investor exited in t (probit-only design with no Stage 2, since the allocation share is mechanically undefined for exiters); Panel C also restricts to held-in-$t-1$ investor-months and sets $D = 1$ if the investor continued in t, 0 if exited. Each panel reports four return specifications. Sub-blocks (.1) and (.2) use the positive/negative decomposition: $r^+ = \max(r, 0)$ and $r^- = \max(-r, 0)$ entered jointly, after winsorising at 1/99 and z-scoring within month. Sub-blocks (.3) and (.4) use a high-vs-median indicator: $r_{\text{high}} = \mathbf{1}\{r > \text{within-month median}\}$. Sub-blocks (.1) and (.3) use raw returns; (.2) and (.4) use excess returns over the CSI 300. The return regressors are the investor's mutual fund portfolio return (Fund) and robo-advisor portfolio return (RA). The investor's mutual-fund return is well defined at every at-risk month and is the only return signal available to predict first adoption. Panels B and C condition on prior RA use (at-risk = held RA in $t-1$), so the RA return is well defined and reported. Observations and T at the panel foot are reported for the leading sub-block; the remaining sub-blocks within a panel use samples within 0.1% of the reported values. Baseline controls (log total financial assets$_{t-1}$, stated risk preference, age, female indicator, combined-advisor dummy$_{t-1}$, cash-management assets$_{t-1}$, and province dummies) are included and suppressed. Months-since-registration ($t-1$) enters Stage 1 as the exclusion restriction. A caveat on the Heckman structural diagnostics for Panels A and C: $|\hat{\rho}| > 1$ in the first-adoption design (the within-month residual sd is small relative to the IMR loading) and $\hat{\rho}$ is similarly large in continuation. The reported IMR coefficients should be read as Heckman diagnostics rather than strict correlations. Stage 2 coefficients on baseline controls and return regressors remain consistently estimated Fama–MacBeth averages. Newey–West HAC t-statistics (lag = $\min(3, T-1)$) in parentheses. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed).

Table A.15: Robo-Advisor Dose-Response on the Mutual-Fund Sleeve: Secondary Outcomes and Placebo Tests

	(1) Auto-invest share	(2) Gambling bias	(3) Disposition effect
RA-share bins at $t-1$ (Bin 0: $rsh_{t-1}=0$ omitted as reference)			
Bin 1 (0–0.20]	0.064***	-0.004	0.002
	(9.207)	(-0.867)	(0.458)
Bin 2 (0.20–0.40]	0.054***	-0.000	0.004
	(7.427)	(-0.075)	(0.691)
Bin 3 (0.40–0.60]	0.029***	-0.002	0.001
	(3.165)	(-0.321)	(0.057)
Bin 4 (0.60–0.80]	0.023**	-0.002	0.037***
	(2.096)	(-0.227)	(2.900)
Bin 5 (0.80–1.00]	-0.103***	0.003	0.023
	(-8.075)	(0.262)	(1.413)
log(TFA), $t-1$	0.001	-0.016**	-0.010**
	(0.160)	(-2.181)	(-2.008)
Baby assets, $t-1$	-0.027***	-0.003*	-0.000
	(-13.390)	(-1.744)	(-0.160)
Combined-advisor dummy, $t-1$	0.036***	0.000	0.007
	(5.352)	(0.008)	(1.212)
Investor FE	Yes	Yes	Yes
Month FE	Yes	Yes	Yes
Sample	MF buy active	MF active	MF active
N	2,509,156	4,159,347	4,016,052
Within R^2	0.002	0.000	0.000

Notes. The table reports two-way fixed-effects panel regressions (`PanelOLS` with investor and month effects) of three secondary mutual-fund (MF) sleeve outcomes on lag-1 robo-advisor (RA) intensity, companion to the headline outcomes in the main text. The specification, treatment construction (six lag-1 RA-share bins with Bin 0 omitted), controls, and two-way-clustered inference are identical to the main table; see those notes for full detail. The three outcomes serve distinct roles. Column (1), the auto-invest share of mutual-fund purchases, is estimated on the smaller MF-buy-active subsample (investors who purchased any fund that month) and traces a dose-response that rises at low RA intensity and reverses sharply at the top bin. Columns (2) and (3), gambling bias and the disposition effect, are deep behavioural-bias measures of the self-directed sleeve; their near-zero within-R^2 and almost uniformly insignificant bin coefficients indicate that RA intensity does not move these biases. They function as placebo tests: the structural restructuring of the sleeve documented in the main table (substitution, concentration, turnover, risk) is not accompanied by a shift in these underlying biases, which sharpens the interpretation that the main effects reflect portfolio reorganisation rather than a general change in investor type. Outcomes are winsorised at 1/99 and z-scored within month (the auto-invest share is built from raw inputs and then standardised). Standard errors are two-way clustered on investor and month; clustered t-statistics in parentheses. As in the main table, the one-month treatment lag guards against contemporaneous reverse causation but not joint within-month determination; estimates describe within-investor co-movement, not a causal effect. Stars: *** $p<0.01$, ** $p<0.05$, * $p<0.10$ (two-tailed).

Table A.16: Time-Stability and Cohort-Split Heckman Outcome-Equation Coefficients

Panel A: Outcome = RA share (RA / (RA + MF)) — BASE spec, outcome equation				
Variable	Early period	Late period	Early cohort	Late cohort
Constant	0.429^{***}	0.073^{***}	0.134^{***}	-0.149
	(3.897)	(8.540)	(3.590)	(-0.323)
Log(TFA) (t-1)	-0.216^{***}	-0.088^{***}	-0.107^{***}	-0.021
	(-6.009)	(-43.846)	(-7.686)	(-0.386)
Risk Preference	-0.019^{**}	0.008^{***}	0.010^{***}	0.007
	(-2.261)	(30.720)	(2.740)	(1.551)
Age	-0.015^{***}	-0.002^{***}	-0.007^{***}	-0.006^{***}
	(-4.393)	(-5.510)	(-3.740)	(-3.161)
Female	0.009	-0.011^{***}	-0.013^{***}	-0.024
	(1.270)	(-99.600)	(-4.300)	(-1.008)
Combined RA Dummy (t-1)	-0.039	0.020^{***}	0.012	0.029
	(-1.440)	(5.420)	(1.390)	(0.456)
Baby Assets (t-1)	0.033^{***}	0.027^{***}	0.020^{***}	0.015
	(5.600)	(26.460)	(8.820)	(1.188)
Inverse Mills Ratio	-0.025	0.169^{***}	0.081^{***}	0.317
	(-0.390)	(21.390)	(3.910)	(1.239)
Province dummies	Yes	Yes	Yes	Yes
T (months)	6	6	48	23

Notes to Panel A. Heckman outcome-equation coefficients, BASE specification (demographics only), for four sample splits: (i) *Early period* — first 6 months of the panel; (ii) *Late period* — final 6 months; (iii) *Early cohort* — investors whose first RA adoption occurred in the first 6 months; (iv) *Late cohort* — investors whose first RA adoption occurred in the final 6 months (the full 48-month panel spans the observation window). Fama–MacBeth with Newey–West HAC t-statistics (max lag $= 3$) in parentheses; reported coefficients are averages of the monthly cross-section estimates. Wide standard errors in the Early-period and Late-cohort columns reflect the short time-series on which these averages are computed ($T = 6$ and $T = 23$ respectively). *Headline.* Signs of demographic effects are stable across the four splits: Log(TFA) is negative, Age is negative, Baby Assets is positive, and Female is (weakly) negative everywhere except Early period. The magnitude of the Log(TFA) effect shrinks over time (from -0.216 in Early period to -0.088 in Late period), suggesting that the wealth effect on intensive delegation weakens as RA technology matures. $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.10$.

Panel B: Outcome = Risky RA share — BASE spec, outcome equation				
Variable	Early period	Late period	Early cohort	Late cohort
Log(TFA) (t-1)	-0.363^{***}	-0.136^{***}	-0.157^{***}	-0.069
	(-6.717)	(-36.830)	(-7.527)	(-1.162)
Risk Preference	-0.059^{***}	0.005^{***}	0.003	-0.001
	(-4.118)	(11.896)	(0.532)	(-0.158)
Age	0.006	-0.007^{***}	-0.011^{***}	-0.008^{***}
	(0.794)	(-21.477)	(-5.265)	(-3.366)
Female	0.050^{***}	-0.002^{***}	-0.004	0.002
	(3.799)	(-8.082)	(-0.867)	(0.060)
Combined RA Dummy (t-1)	-0.114^{***}	-0.017^{***}	-0.019^{*}	-0.032
	(-3.586)	(-3.348)	(-1.772)	(-0.366)
Baby Assets (t-1)	0.061^{***}	0.040^{***}	0.030^{***}	0.028^{*}
	(7.001)	(29.698)	(9.605)	(1.834)
Province dummies	Yes	Yes	Yes	Yes

Notes to Panel B. Same four sample splits as Panel A but using the *risky* RA share — ra_risky/(ra_risky + risky_mf) — as the outcome. This is the within-risky-portfolio delegation share, holding the risk allocation fixed. Fama–MacBeth with Newey–West HAC t-statistics (max lag = 3) in parentheses. *Headline.* Sign patterns again stable across splits. Log(TFA) is strongly negative everywhere (wealthier investors tilt less toward risky RA). Baby Assets and Risk Preference (late period) are positive. The Female coefficient flips sign between Early period ($+0.050^{***}$) and Late period (-0.002^{***}), possibly reflecting rapid adoption by female investors in the early window when RA was a novel product. Constants and Inverse Mills ratios are reported in Panel A's footnote only to save space. $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.10$.

Table A.17: Correlation of Age with Behavioral Biases, Trading Skills, and Experience

	Investor-Month Panel			Investor-Level (Means)		
Variable	N	Pearson	Spearman	N	Pearson	Spearman
Panel A: Anchoring Biases						
Anchoring Bias	4,317,961	$+0.0214^{***}$	$+0.0204^{***}$	100,000	$+0.0404^{***}$	$+0.0374^{***}$
Anchoring Buy Ratio	—	—	—	100,000	$+0.0183^{***}$	$+0.0134^{***}$
Anchoring Sell Ratio	4,085,758	$+0.0168^{***}$	$+0.0113^{***}$	100,000	$+0.0299^{***}$	$+0.0223^{***}$
Panel B: Turnover						
Total Turnover	4,402,068	-0.0675^{***}	-0.1097^{***}	98,071	-0.0970^{***}	-0.1484^{***}
Purchase Turnover	4,402,068	-0.0256^{***}	-0.0394^{***}	98,071	-0.0479^{***}	-0.1294^{***}
Sale Turnover	4,402,068	-0.0631^{***}	-0.0813^{***}	98,071	-0.0957^{***}	-0.1326^{***}
Speculative Turnover	4,356,528	$+0.0036^{***}$	-0.0034^{***}	98,071	$+0.0057$	-0.0507^{***}
Spec. Purchase Turnover	4,302,707	-0.0039^{***}	$+0.0052^{***}$	98,071	-0.0104^{**}	-0.0579^{***}
Spec. Sale Turnover	—	—	—	98,028	$+0.0185^{***}$	-0.0428^{***}
Speculative Turnover (Type)	4,356,528	-0.0151^{***}	-0.0108^{***}	98,071	-0.0275^{***}	-0.0705^{***}
Panel C: Attention						
Attention Bias	4,670,815	-0.0258^{***}	-0.0303^{***}	100,000	-0.0608^{***}	-0.0691^{***}
Panel D: Disposition Effect						
PGR	4,565,152	-0.0620^{***}	-0.0509^{***}	99,986	-0.1093^{***}	-0.1462^{***}
PLR	4,441,930	-0.0632^{***}	-0.0534^{***}	99,720	-0.1041^{***}	-0.1530^{***}
Disposition Effect (DE)	—	—	—	98,342	-0.0138^{***}	-0.0149^{***}
DE Ratio	3,198,314	$+0.0179^{***}$	-0.0030^{***}	85,233	$+0.0163^{***}$	-0.0146^{***}
PGR (value-weighted)	4,565,152	-0.0569^{***}	-0.0502^{***}	99,986	-0.1039^{***}	-0.1415^{***}
PLR (value-weighted)	4,441,930	-0.0597^{***}	-0.0528^{***}	99,720	-0.1010^{***}	-0.1508^{***}
DE (value-weighted)	4,268,836	-0.0059^{***}	-0.0023^{***}	98,342	-0.0114^{***}	-0.0110^{***}
Panel E: Diversification & Concentration						
Underdiv. HHI	4,717,162	-0.0344^{***}	-0.0529^{***}	100,000	-0.0417^{***}	-0.0661^{***}
Portfolio HHI	4,717,162	-0.0344^{***}	-0.0529^{***}	100,000	-0.0417^{***}	-0.0661^{***}
Underdiv. (Corr.)	3,568,973	$+0.0165^{***}$	$+0.0299^{***}$	97,688	$+0.0195^{***}$	$+0.0384^{***}$
Underdiv. (Div.)	3,570,516	$+0.0074^{***}$	$+0.0137^{***}$	97,699	$+0.0121^{***}$	$+0.0220^{***}$
# Funds Held	4,717,162	$+0.0884^{***}$	$+0.0608^{***}$	100,000	$+0.0953^{***}$	$+0.0550^{***}$
Effective # Funds	4,717,162	$+0.0699^{***}$	$+0.0516^{***}$	100,000	$+0.0789^{***}$	$+0.0462^{***}$
Diversification Ratio	3,570,516	-0.0064^{***}	-0.0134^{***}	97,699	-0.0106^{***}	-0.0280^{***}
Correlation-Weighted Conc.	3,568,973	$+0.0165^{***}$	$+0.0299^{***}$	97,688	$+0.0195^{***}$	$+0.0384^{***}$
Panel F: Extrapolation						
Extrapolation Bias	4,612,051	-0.0163^{***}	-0.0214^{***}	100,000	-0.0430^{***}	-0.0533^{***}
Extrapolation Bias (3m)	4,545,865	-0.0228^{***}	-0.0289^{***}	100,000	-0.0543^{***}	-0.0732^{***}
Panel G: Gambling						
Gambling Bias	4,475,220	-0.0173^{***}	-0.0150^{***}	100,000	-0.0371^{***}	-0.0460^{***}
Gambling Bias (Strict)	4,213,337	-0.0037^{***}	$+0.0042^{***}$	100,000	-0.0095^{**}	-0.0198^{***}
Panel H: Confirmation						
Confirmation Bias	3,520,292	$+0.0603^{***}$	$+0.0572^{***}$	99,705	$+0.0981^{***}$	$+0.1232^{***}$
Disconfirmation Bias	3,319,970	-0.0687^{***}	-0.0665^{***}	98,235	-0.1021^{***}	-0.1465^{***}
Panel I: Herding						
Herding (Buy)	4,596,186	-0.0307^{***}	-0.0349^{***}	100,000	-0.0542^{***}	-0.0613^{***}
Herding (Sell)	4,273,876	-0.0243^{***}	-0.0173^{***}	100,000	-0.0401^{***}	-0.0052
Herding Bias	4,612,051	-0.0265^{***}	-0.0299^{***}	100,000	-0.0497^{***}	-0.0474^{***}
Panel J: Skills & Experience						

(Continued)

	Investor-Month Panel			Investor-Level (Means)		
Variable	N	Pearson	Spearman	N	Pearson	Spearman
Selection Skill (S_t)	—	—	—	99,898	$+0.0357^{***}$	$+0.0486^{***}$
Timing Rate	4,571,144	$+0.0005$	-0.0045^{***}	99,891	$+0.0014$	$+0.0026$
Fund Experience (months)	4,645,675	$+0.0807^{***}$	$+0.1180^{***}$	100,000	$+0.0835^{***}$	$+0.1417^{***}$
RA Experience (months)	2,377,855	$+0.0401^{***}$	$+0.0571^{***}$	70,049	$+0.0311^{***}$	$+0.0518^{***}$

Notes. Pearson and Spearman rank correlations between investor age and each behavioral, skill, or experience measure. The investor-month panel uses the winsorised (1%/99%) and within-month standardised series; the investor-level columns collapse the panel to one observation per investor by averaging across months. Age is in years (raw, not standardised). N is the number of pairwise non-missing observations. Significance reflects the test that the population correlation is zero: $^{*}\ p < 0.05$, $^{**}\ p < 0.01$, $^{***}\ p < 0.001$. In the investor-month panel the very large N (≈2.4–4.7 million) drives most p-values below 10^{-50}; the magnitude of the correlation, not statistical significance, is the relevant economic statistic. Investor-level correlations are systematically larger in absolute value (typically by a factor of 2–3) because aggregating to the investor mean removes within-investor month-to-month noise. Variable definitions follow Table 1.

Table A.18: The Structure of the Self-Directed Sleeve Within the RA-Adopter Sample

Sample: adopter-month panel (investors who hold RA in at least one sample month). *Estimator:* Fama–MacBeth OLS.
Dependent variables measured contemporaneously at month t.

Regressor	(1) # Funds Held	(2) Portfolio HHI	(3) Self-Directed Share	(4) Total Turnover	(5) Speculative Turnover
Constant	10.281***	0.403***	0.787***	0.283***	0.014***
	(34.167)	(37.252)	(65.420)	(17.078)	(8.136)
Log(TFA)$_{t-1}$	7.117***	-0.158***	0.114***	-0.147***	-0.003***
	(40.125)	(-32.466)	(14.882)	(-14.298)	(-8.374)
Risk Preference	0.759***	-0.003***	-0.010***	0.011***	0.000*
	(13.925)	(-3.525)	(-7.324)	(6.558)	(1.920)
Age	0.755***	0.011***	0.007***	0.005***	0.001***
	(29.920)	(11.470)	(5.249)	(5.809)	(7.522)
Female	-1.247***	-0.009***	0.004***	-0.035***	-0.003***
	(-21.829)	(-12.534)	(5.151)	(-20.842)	(-10.202)
Combined RA Dummy$_{t-1}$	8.927***	-0.065***	0.022***	-0.043***	-0.004***
	(18.519)	(-10.890)	(11.185)	(-12.959)	(-8.133)
Baby Assets$_{t-1}$	-1.773***	0.031***	-0.023***	0.042***	0.002***
	(-31.001)	(41.477)	(-14.432)	(31.614)	(5.484)
Fund Experience (std)$_{t-1}$	2.158***	-0.051***	0.041***	-0.065***	-0.001*
	(27.856)	(-15.822)	(9.171)	(-14.957)	(-1.740)
Selection Skill (std)$_{t-1}$	-0.061	0.002	0.002**	0.004	0.001**
	(-1.020)	(1.373)	(2.001)	(1.602)	(2.393)
Timing Rate (std)$_{t-1}$	0.010	-0.000	0.001**	-0.000	0.001***
	(0.266)	(-0.600)	(2.062)	(-0.058)	(3.214)
Province Dummies	Yes	Yes	Yes	Yes	Yes
N	1,530,328	1,530,328	1,530,328	1,485,248	1,485,248
Avg. R^2	0.152	0.206	0.145	0.142	0.009

Notes. Fama–MacBeth OLS on the *adopter-month panel* (investors who hold RA in at least one sample month); no selection correction. Each month, a cross-sectional OLS is estimated on the adopter panel; reported coefficients are the time-series mean of the monthly betas, with Newey–West HAC t-statistics (max lag $= 3$) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed). *Dependent variables* at month t: (1) # Funds Held — count of distinct mutual funds in the self-directed book; (2) Portfolio HHI — Herfindahl concentration of self-directed fund weights; (3) Self-Directed Share $=$ MF assets / (RA $+$ MF); (4) Total Turnover — (buys $+$ sells)/lagged holdings in the self-directed book; (5) Speculative Turnover — turnover in speculative/short-horizon positions. *Controls* (all lagged one month): Log(Total Financial Assets), Combined RA Dummy, Baby Assets, Fund Experience, Selection Skill, Timing Rate. Age, Female, and Risk Preference are time-invariant and not lagged. Selection Skill and Timing Rate are standardised within month. Province dummies are included but not reported. *Key finding.* The coefficient on the lagged Combined RA Dummy describes the self-directed sleeve of an adopter in months with RA activity relative to adopter-months without it. In these months, adopters hold more self-directed funds ($+8.93$), exhibit lower portfolio concentration (-0.065), carry a larger self-directed share ($+0.022$), and trade less in both overall (-0.043) and speculative (-0.004) terms. The evidence is consistent with RA use coexisting with a broader, less concentrated, and more disciplined self-directed portfolio, rather than serving as a direct substitute for self-directed investing.