The Demand for Robo-Advising: Adoption and Delegation Intensity

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

May 31, 2026

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

Abstract

Robo-advisors provide scalable financial assistance to retail investors through delegated portfolio management, yet evidence on who adopts this form of assistance remains conflicting. We argue that these conflicting results reflect a lack of distinction between adoption and delegation intensity conditional on adoption. Using proprietary monthly Alipay data on 100,000 Chinese mutual fund investors from September 2020 through September 2024, we employ the Heckman selection framework to disentangle the two margins. We find systematic evidence of opposite signs: the same investor characteristic predicts adoption in one direction and delegated share in the other. Underdiversification and turnover reduce adoption but increase adopters' delegated share. By contrast, confirmation bias and fund experience increase participation, but reduce delegation. To explain the observed patterns of adoption and delegation, we propose and test three mechanisms: selective awareness, decision authority, and reclaiming control. First, selective awareness explains why investors with correctable mistakes are less likely to adopt robo-advising, but delegate more and improve portfolio behavior after exposure. Second, decision authority distinguishes robo-advising from generic automation: within-investor-month comparisons show that automating execution and delegating fund-selection authority are distinct choices. Third, reclaiming control explains why more experienced investors reduce delegation over time. The paper contributes to the behavioural and household finance literature by identifying the determinants of selection into investment assistance technologies and showing that adoption and delegation are distinct dimensions of household financial decision-making. The results reveal how behavioural biases shape both the demand for investment assistance and the extent to which investors are willing to cede decision authority.

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

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

*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 desensitised, and were remotely analysed on the Ant Open Research Laboratory in an Ant Group environment. We thank Chris Brooks, Lorenz Kueng, 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.

1. Introduction

Algorithms now mediate a growing share of household financial decisions, with consequences for diversification, trading, and long-run wealth accumulation. Robo-advisors are a leading case because they make delegated portfolio management available to retail investors at scale. Their economic role rests on two choices rather than one. An investor must first decide whether to adopt the service, and then decide how much wealth, and how much fund-selection authority, to delegate after adopting.

Existing evidence documents both who adopts robo-advisors and who benefits from them, yet it remains unclear how investors demand delegated investment management and whether adoption and delegation are distinct choices. Some studies find that robo-advisor adopters are younger, male, and hold smaller portfolios (Baulkaran & Jain, 2023; Brenner & Meyll, 2020; Hao et al., 2022); others find weak or insignificant age and gender effects (Belanche et al., 2019; D'Acunto et al., 2019). The behavioural finance literature finds that overconfidence raises adoption while measured financial literacy depresses it (Piehlmaier, 2022). At the same time, studies of robo-advisor outcomes show that use can improve diversification, reduce mistakes, and raise portfolio quality (Bianchi & Brière, 2026; Hao et al., 2022; Reher & Sokolinski, 2024; Rossi & Utkus, 2024). The tension is that evidence on who adopts the technology does not map cleanly onto evidence on who appears to benefit from it. We argue that this conflict is not primarily a sample or measurement problem, but a specification problem, because the literature has usually treated use as a binary decision, when the demand for delegated portfolio management has both an entry margin and a commitment margin. This paper therefore asks whether that demand is a single participation decision or two choices governed by distinct economic primitives, a fixed cost of entry and a marginal cost of ceding decision authority, and whether the same investor characteristic that raises adoption can lower the share of wealth and fund-selection authority delegated once inside.

The distinction is economically important because the two margins reflect different costs. Entry requires discovery, evaluation, and initial trust. Commitment requires the investor to surrender additional fund-selection and fund-trading authority. The same investor characteristic can affect these costs in opposite ways. An underdiversified investor may be less likely to adopt because she does not recognize concentration as costly, yet delegate more once she adopts and sees a diversified algorithmic portfolio. An experienced fund investor may know enough to try the product, but delegate little because fund selection and trading are the decisions she values. The same characteristic can therefore inc adoption and reduce delegated share, or vice versa. A binary adoption variable cannot separate these cases, and any literature built on it will recover whichever margin happens to dominate in a given sample.

Distinguishing between adoption and delegation in robo-advising has been difficult because the required evidence is rarely observed in a single setting. Surveys measure willingness to use robo-advice but not realized delegated shares after adoption. Platform data record product use, but often lack information on the self-managed investments from which investors are delegating, making delegation difficult to quantify. Portfolio datasets observe holdings but rarely allow comparison between delegated investing and execution-only automation within the same account. Cross-sectional data often combine first adoption, continued use, exit, re-entry, and later reallocation. Testing whether

adoption and delegation are governed by different economic forces requires observing, for the same investor over time, adoption, delegated share, self-directed assets, trades, behavioral characteristics, and a comparable product that automates execution without transferring fund-selection authority. We study BangNiTou, a fully delegated mutual fund robo-advisor founded by Ant Group and Vanguard and embedded within Alipay. The same Alipay account records self-directed mutual fund investments, robo-advised assets, automatic investment plans, cash-management balances, platform engagement, transaction histories, and realized returns. The monthly panel of 100,000 Chinese mutual fund investors from September 2020 through September 2024, totaling 4.9 million investor-months, allows us to study entry, active use, tenure, and post-adoption reallocation. The platform also offers automatic investment plans, providing a useful comparison with a form of investment automation that does not delegate fund-selection authority.

Because delegated share is observed only for adopters, we estimate monthly Heckman two-stage selection models and report Fama-MacBeth averages with Newey-West standard errors. The first stage estimates adoption; the second stage estimates the robo-advised share of the active fund portfolio among adopters, correcting for selection into adoption. The exclusion variable is platform tenure, which is intended to shift product discovery and evaluation while not directly determining the delegated share once adoption and controls are held fixed. The design gives empirical content to a simple economic distinction: trying the product reflects a fixed entry friction, whereas increasing the delegated share reflects a marginal cost of ceding authority.

The estimates show that the two margins are not governed by the same investor characteristics, and that several characteristics carry opposite signs across them. Wealth primarily affects the adoption margin. In the baseline specification, a one-standard-deviation increase in log total financial assets (2.6 log points) raises adoption by 14.1 percentage points, about 43 percent of the sample participation rate. Its negative association with delegated share, initially about half of the 0.151 conditional mean, falls to near zero and loses significance once we control for measures of behavioral biases. Age becomes a positive adoption predictor after biases are controlled for, while its delegated-share relation vanishes. Fund experience raises the probability of adoption by about 10 percent of the participation mean but lowers conditional delegation by about 19 percent of the conditional mean, and it predicts lower delegated shares at longer robo-advisor tenure. Experience therefore plays a dual role: it can lower search and evaluation frictions while raising the value of retaining fund-selection authority.

Because robo-advising might help mitigate behavioral trading patterns through diversified, rule-based portfolio construction, an important question is whether the investors who stand to benefit most are also the ones who adopt it. We therefore next study behavioral patterns in investing. Underdiversification and turnover reduce adoption but increase delegated share among adopters. A one-standard-deviation increase in portfolio concentration, lowers adoption by 6.6 percentage points, or about 18 percent of the participation mean, and raises the conditional share by 2.0 percentage points, or about 13 percent of the conditional mean. Turnover lowers adoption by about 8 percent of the participation mean and raises the conditional share by about 27 percent of the conditional mean, indicating a much stronger effect on delegation intensity than on participation. Confirmation bias shows the reverse pattern, raising adoption but reducing delegated share. Attention bias, gambling preference, and herding reduce both margins. Behavioral factors are, therefore, central determinants of demand. Rather than

affecting participation uniformly, they operate through distinct margins: some discourage adoption while increasing delegation among adopters, some limit commitment after entry, and some reduce both margins.

Overall, the findings show that investors are least likely to adopt robo-advice when they exhibit the very behaviors that the technology is most effective at addressing (i.e., underdiversification and high turnover). Yet once these investors adopt, they delegate more than other adopters. Post-adoption evidence reinforces this interpretation. Among adopters, underdiversification and turnover decline over robo-advisor tenure, with larger declines among initially high-bias and more engaged investors. We refer to this mechanism as selective awareness. Adoption reveals the value of delegation only for mistakes that the algorithm visibly corrects. Investors who delegate observe a diversified, rule-based portfolio reducing concentration and discretionary trading, and update their beliefs about the benefits of delegation accordingly. By contrast, delegation does not directly relax cognitive constraints such as limited attention, gambling preference, or herding. Because these behaviors are not visibly corrected through exposure, they continue to depress both adoption and delegation. Awareness is therefore selective across mistakes rather than uniform, which explains why underdiversification and turnover reverse sign across the two margins while cognitive-constraint biases do not. A further test examines how investors come to recognize the value of delegation. One possibility is a return-chasing channel: investors may delegate more after observing that robo-advised portfolios outperform their self-directed investments. If so, the lagged market-versus-robo return gap should strengthen the relationship between underdiversification and delegation. We find no such effect. Instead, the evidence is more consistent with selective awareness, under which investors learn from the changes delegation produces within their own portfolios. The cue is not relative performance but the visible reduction in concentration and discretionary trading that follows adoption.

The next test asks whether the key mechanism for adoption is financial automation in general or the delegation of decision authority in particular. We compare three choices within the same investor-month: self-directed investing, automatic investment plans, and robo-advising. Automatic investment plans automate execution while preserving the investor's fund choice; robo-advising also delegates fund selection. Investor-month fixed effects absorb wealth, market conditions, and any general taste for automation, so the comparison isolates the effect of delegating fund-selection authority. Skill and engagement do not reflect a common demand for automation. More skilled investors are more likely to adopt execution automation but are no more likely to delegate fund-selection authority, while more engaged investors are less likely to automate execution yet more likely to delegate investment decisions. These contrasting patterns are difficult to reconcile with a single taste for automation.

Because a static adoption model cannot tell whether delegation is an absorbing state, we study how delegated share evolves with robo-advisor tenure and how this varies with pre-adoption fund experience. Among active adopters, longer tenure is associated with lower delegated shares, and the decline is larger for investors with more pre-adoption fund experience. Experienced investors do not simply adopt and remain delegated; they hold lower shares at longer tenures, consistent with reclaiming control as their self-directed capability raises the cost of ceding authority. Event-study evidence around adoption shows the same pattern on the self-directed side: portfolio concentration falls sharply at adoption and then recovers toward its pre-adoption level over the following months.

Delegation, therefore, coincides with a temporary, not a permanent, reshaping of the investor's self-directed portfolio.

Two further results add important nuance to the interpretation. Selection on unobservables is positive and economically large. The inverse Mills coefficient rises across specifications, from 0.186 in the baseline panel to 0.251 with behavioral controls, and remains stable after adding a rich set of skill and bias controls that might have absorbed it. Observably similar adopters therefore differ on an unobserved dimension that raises conditional delegation, which we interpret as latent trust or comfort with delegated algorithmic management rather than as any single measured trait, in the spirit of evidence that trust predicts financial participation beyond wealth and sophistication (Guiso et al., 2008). Prior investment performance affects some adoption, continuation, and exit decisions in some specifications, but it does not account for the contrasting effects of investor characteristics on adoption and delegation. The magnitude of recent own-portfolio returns, in either direction, lowers the probability of first adoption, while at exit and continuation, it is the sign of returns that drives the decision. Prior performance affects lifecycle margins without explaining the central cross-margin patterns in the data. Post-adoption evidence then shows that entry and delegated share are associated with different adjustments in the self-directed portfolio. While largely descriptive, these results complement the core evidence and point to a demand system shaped by fixed entry frictions, marginal authority costs, selective awareness of mistakes that delegation can address, latent trust, and dynamic reallocation between delegated and self-directed investing.

The paper makes at least four contributions to the household and behavioral finance literature. First, separating adoption from conditional delegation, we reconcile the contrasting findings in the literature concerning who uses robo-advice (Baulkaran & Jain, 2023; Belanche et al., 2019; Brenner & Meyll, 2020; D'Acunto et al., 2019; Hao et al., 2022; Piehlmaier, 2022). The conflicting findings are largely a specification artifact, since a binary use indicator pools entry with commitment, so samples that differ in composition recover different mixtures of two margins on which the same characteristic can carry opposite signs. The same separation also sheds light on the relationship between participation and portfolio quality in household finance (Calvet et al., 2007; Campbell, 2006). For delegated investment products, an additional margin lies between participation and outcomes: the extent to which households delegate wealth and decision authority. Consequently, greater participation need not translate into greater portfolio improvement.

Second, we treat behavioral investment patterns as the object that governs demand for a corrective technology rather than as nuisance controls. Behavioral finance documents costly household investment mistakes (Barber & Odean, 2000, 2001, 2008; Goetzmann & Kumar, 2008; Kumar, 2009), and the robo-advice effects literature shows that algorithms can reduce them conditional on use (Bianchi & Brière, 2026; Hao et al., 2022; Reher & Sokolinski, 2024; Rossi & Utkus, 2024). Both take use as given and study either the mistake or its improvement; neither study selection into the correction. We add the missing demand side of this process by studying which investors choose to adopt robo-advice and how much decision authority they delegate once they do. This distinction matters because the aggregate gains from corrective technology depend not only on its effectiveness but also on whether the investors who stand to benefit most are willing to expose themselves to it.

Third, we separate the delegation of decision authority from financial automation. The literature

on delegated management and algorithm aversion (Dietvorst et al., 2015, 2018; Guiso et al., 2008) typically bundle automating a task with transferring a decision, so willingness to automate cannot be distinguished from willingness to delegate. Our within-investor-month comparison across self-direction, execution automation, and full delegation isolates the authority transfer in observed portfolio decisions. The distinction extends to managed accounts, target-date funds, and other products that automate tasks while varying in the degree of decision authority they transfer.

Fourth, we show that delegation is dynamic rather than static. Adoption studies treat use as an absorbing state; we document a partly reversible post-adoption adjustment in which more experienced investors reclaim control over tenure. This links the delegation decision to learning-by-trading in household finance (Linnainmaa, 2011; Nicolosi et al., 2009; Seru et al., 2010) and shows that the degree of control investors cede continues to evolve after adoption.

The rest of the paper proceeds as follows. Section 2 presents the conceptual framework. Section 3 describes the data and empirical designs. Section 4 reports the baseline Heckman evidence and the mechanism tests it motivates. Section 5 studies first adoption and lifecycle margins. Section 6 reports descriptive evidence on consequences for the self-directed portfolio. Section 7 synthesizes the framework and results. Section 8 summarizes robustness. Section 9 concludes.

2. Conceptual Framework

This section develops a reduced-form framework for interpreting demand for robo-advising. The framework is deliberately parsimonious. It separates the margins that the empirical analysis can observe. The central object is a participation decision that requires delegation. This distinguishes robo-advising from standard market entry in household finance, where participation and portfolio efficiency often move in the same direction (Calvet et al., 2007; Campbell, 2006). In a delegated product, participation requires the investor to give up fund-selection authority, so the extensive and intensive margins need not have the same sign.

The framework serves three purposes. First, it explains why the same investor characteristic can have opposite signs on adoption and delegated share. Second, it classifies behavioral variables by the economic reason their signs differ across margins. Third, it justifies the auxiliary tests used later in the paper: the DIRECT-AUTO-RA authority ladder, post-adoption bias trajectories, performance-prediction exercises, and tenure-by-experience dynamics. The framework is therefore a map for the empirical design rather than a structural model to be estimated.

2.1. Environment and Timing

Investor i is observed in month t with platform fund wealth normalized to one. Three product modes are available within the same account. DIRECT denotes self-directed mutual fund investing. AUTO denotes automatic investment plans, which automate purchase timing but leave fund selection to the investor. RA denotes robo-advising, which automates execution and transfers fund-selection and fund-trading authority to an algorithm.

The timing is as follows. First, the investor decides whether to use RA. Second, conditional on RA use, she chooses the share $s_{it} \in [0, 1]$ of the RA-plus-mutual-fund wallet to allocate to RA. Third,

after adoption, the investor observes experience with the RA sleeve and may adjust both the delegated share and behavior in the self-directed sleeve. This timing maps to the empirical sections: the baseline Heckman estimates adoption and conditional share; the authority-ladder design separates DIRECT, AUTO, and RA; the post-adoption tests study bias trajectories and RA tenure.

The key modelling choice is to treat adoption and delegated share as distinct choices rather than as a single latent demand index. A one-index model can generate different magnitudes across margins, but it has difficulty generating stable opposite signs without auxiliary assumptions. The two-cost structure below makes the opposite signs natural: one cost is paid once at entry, while the other is paid on every unit of wealth delegated to the algorithm.

2.2. Fixed Entry Costs and Marginal Authority Costs

The investor adopts RA when the indirect value of access exceeds a fixed entry cost:

$$d_{it} = 1 \quad \Longleftrightarrow \quad V_{it}^{RA} - F_{it} + \nu_{it} \geq 0, \tag{1}$$

where F_{it} combines search, product discovery, onboarding, and initial trust formation, and ν_{it} is an unobserved taste or trust component. Let

$$F_{it} = F_0 - \alpha_W W_{it} - \alpha_R R_{it} - \alpha_P P_{it} + \alpha_A A_{it} + \varepsilon_{it}^F. \tag{2}$$

Here W_{it} is platform wealth, R_{it} is prior robo-product exposure, P_{it} is platform familiarity, and A_{it} is an attention or evaluation constraint. Wealth and prior platform experience can lower the proportional cost of trial. Attention constraints raise the cost of evaluating a new financial product.

Equation (2) also motivates the Heckman exclusion restriction. Platform familiarity is allowed to shift discovery and evaluation costs. Conditional on adoption and the outcome covariates, the identifying assumption is that this discovery channel does not directly alter the optimal share delegated to RA. The assumption is not implied by the model; it is a maintained empirical restriction. The model clarifies its economic content.

Conditional on RA use, the delegated share solves a separate allocation problem. A reduced-form representation is

$$s_{it}^* = \left[\frac{\Delta r_{it} + B_{it}^{RA} - C_{it}^{auth}}{\kappa_i} \right]_0^1, \tag{3}$$

where Δr_{it} is the perceived service or return advantage of RA over the self-directed sleeve, B_{it}^{RA} is the perceived benefit of algorithmic discipline, C_{it}^{auth} is the marginal cost of giving up fund-selection authority, $\kappa_i > 0$ captures curvature, and $[\cdot]_0^1$ truncates the interior solution to the unit interval. The marginal authority cost is

$$C_{it}^{auth} = \gamma_E E_{it} + \gamma_S S_{it} + \gamma_B B_{it}^{keep} + \gamma_A A_{it} + \varepsilon_{it}^M, \tag{4}$$

where E_{it} is fund-investing experience, S_{it} is selection or timing skill, and B_{it}^{keep} captures behavior that reveals attachment to self-directed authority, such as concentrated active positions or speculative round trips.

The marginal cost term should be read broadly. It includes explicit fees, the loss of discretion over fund selection, psychological costs of giving up control, and the opportunity cost of not using one's own fund-market human capital. The empirical tests do not separately identify these components. The authority-ladder design is useful because it holds execution automation fixed when moving from AUTO to RA and varies the authority component most directly.

One way to microfound equation (3) is to write the incremental payoff from delegating share s as

$$U_{it}^{RA}(s) - U_{it}^{Self}(s) = \left(\Delta r_{it} + B_{it}^{RA} - C_{it}^{auth}\right) s - \frac{\kappa_i}{2} s^2. \tag{5}$$

The first term is the perceived marginal gain from RA relative to the self-directed sleeve. It combines perceived return or service advantage, algorithmic discipline, and the authority cost. The second term captures curvature in the share decision: the marginal psychological and portfolio cost of delegating rises as more of the fund wallet is ceded to the algorithm. Maximizing equation (5) over $s \in [0, 1]$ gives equation (3). The interior derivative is

$$\frac{\partial s_{it}^*}{\partial X} = \frac{1}{\kappa_i} \left(\frac{\partial \Delta r_{it}}{\partial X} + \frac{\partial B_{it}^{RA}}{\partial X} - \frac{\partial C_{it}^{auth}}{\partial X} \right), \tag{6}$$

whenever the solution is not at the boundary. The expression makes clear why the intensive-margin sign need not match the adoption sign. A variable can make RA easier to discover and evaluate through F_{it} while lowering the optimal share through C_{it}^{auth}.

The model is intentionally reduced form about beliefs. The perceived return advantage Δr_{it} need not equal realized recent performance. It may combine beliefs about diversification, expected service quality, trust in the algorithm, and observed returns. The empirical return tests later in the paper therefore test whether realized performance proxies for this perceived advantage; they do not exhaust the contents of Δr_{it}.

Equations (2) and (4) generate the first empirical implication. A characteristic can lower fixed entry costs while raising marginal authority costs. Such a characteristic should increase adoption and reduce conditional delegation. This is the logic behind the opposite-sign interpretation of fund experience and related sophistication proxies.

Prediction 1 (Two-margin sign reversal)**.** *Let X lower the fixed cost of trying RA and raise the marginal authority cost of delegated wealth. Then X can increase* $\Pr(d_{it} = 1)$ *while reducing s_{it}^* conditional on $d_{it} = 1$.*

The prediction is not a welfare claim. It states that adoption and delegated share are different demand margins. A positive adoption coefficient does not imply a positive intensive-margin coefficient.

The prediction also implies a benchmark. If a variable only lowers fixed entry costs and has no authority-cost component, it should affect adoption but not conditional delegation. Wealth and age after behavioral controls are interpreted against this benchmark. If a variable only raises marginal authority costs and does not affect discovery, it should lower conditional delegation with little entry effect. The empirical patterns need not fit either benchmark exactly, but the benchmarks discipline the interpretation.

2.3. Biases and Selective Awareness

Let B_{it} denote the intensity of a correctable active-opinion bias, with underdiversification and excessive turnover as the leading examples. The true cost of the bias is increasing in B_{it}, but the investor may not perceive that cost fully. We write perceived pre-adoption bias cost as

$$\Lambda_{it}^{pre} B_{it}, \qquad \Lambda_{it}^{pre} = \lambda_0 - \lambda_1 B_{it}, \qquad \lambda_1 > 0. \tag{7}$$

The assumption captures selective awareness: investors with more severe correctable mistakes may be less likely to recognize them as costly. This regularity is consistent with the household-finance evidence on concentrated portfolios and overtrading (Barber & Odean, 2000, 2001; Goetzmann & Kumar, 2008; Linnainmaa, 2011). If selective unawareness is strong, B_{it} can reduce adoption even though RA would correct the behavior.

The important restriction is not that biased investors are irrational in every domain. It is that the perceived private value of correction rises less than the true cost of the mistake as the mistake becomes more severe. This is why the demand-side prediction differs from the effects-side prediction. A bias can increase the potential gain from RA while reducing voluntary entry into RA.

Post-adoption exposure changes the relevant margin. Once the investor has used RA, the comparison between the delegated sleeve and the self-directed sleeve can make own-portfolio concentration or trading intensity more salient. Let τ be months since first RA adoption. For correctable active-opinion biases, the framework predicts

$$\frac{\partial B_{it}}{\partial \tau} < 0, \qquad \frac{\partial^2 B_{it}}{\partial \tau \partial H_i} < 0, \qquad \frac{\partial^2 B_{it}}{\partial \tau \partial G_{it}} < 0, \tag{8}$$

where H_i marks high pre-adoption bias and G_{it} captures engagement. The decline should be stronger among investors with more severe pre-adoption mistakes and among investors who are more likely to observe and process the RA comparison.

Prediction 2 (Selective-awareness signature)**.** *A correctable active-opinion bias can have three linked implications: it can reduce adoption, raise conditional delegation among adopters, and decline over RA tenure after adoption. The post-adoption decline should be strongest for high-pre-adoption-bias and more engaged adopters.*

The prediction is qualitative. It does not require every behavioral variable to decline after RA adoption. It applies to biases for which RA changes the decision environment directly by replacing discretionary fund selection and trading with rule-based delegated management.

This prediction motivates the bias-trajectory table. The mechanism label is selective awareness because the demand-side prediction concerns recognition of correctable mistakes before adoption. The post-adoption decline is the dynamic learning-after-exposure implication. The strongest test is therefore not merely a negative average tenure slope, which a pure time trend or mean reversion could generate. The stronger implication is heterogeneity: the decline should be concentrated among high-pre-adoption-bias investors and should be amplified by engagement, because engaged investors are more likely to observe and process the comparison between their self-directed sleeve and the delegated sleeve.

2.4. Attachment and Cognitive Constraints

Other biases have different economic content. Confirmation bias can make trial easier if RA can be added as a separate sleeve without forcing the investor to unwind validated positions. But the same attachment can reduce conditional delegation because a large RA share requires reallocating away from positions or narratives the investor wants to preserve. This produces the reverse opposite-sign pattern: positive adoption and negative conditional delegation.

Formally, attachment can be represented as a variable Q_{it} that lowers the psychological cost of trying a supplementary RA sleeve but raises C_{it}^{auth} once a larger share must be delegated:

$$\frac{\partial F_{it}}{\partial Q_{it}} < 0, \qquad \frac{\partial C_{it}^{auth}}{\partial Q_{it}} > 0. \tag{9}$$

The model does not distinguish attachment to existing winning positions from attachment to a validated self-image. Both mechanisms have the same cross-margin signature, so the empirical interpretation remains deliberately qualified.

Cognitive constraints and preference-like biases have a third signature. Limited attention, gambling preference, and herding need not be corrected by RA adoption. They can raise evaluation costs and lower the utility of a diversified delegated portfolio. In equations (2) and (4), such constraints enter both margins with the same sign.

Prediction 3 (Three-channel bias taxonomy)**.** *The framework predicts three sign signatures. Correctable active-opinion biases can reduce adoption and raise conditional delegation. Attachment biases can raise adoption and reduce conditional delegation. Cognitive-constraint biases can reduce both margins.*

This taxonomy is useful because it links the demand for corrective technology to the behavioral finance literature on the mistakes the technology can potentially correct. It also imposes a falsification discipline. A uniform learning account would predict reversals for all biases; a no-learning account would predict same-signed effects for all biases. The taxonomy predicts heterogeneity across bias classes.

The taxonomy also explains why the baseline Heckman is not the end of the analysis. The sign pattern points to a candidate mechanism; the post-adoption trajectory evidence asks whether the behavior actually changes after exposure. For attachment and cognitive-constraint biases, the framework does not require the same trajectory test to be positive. Those mechanisms can be supported by their cross-margin signs and by their contrast with the focal active-opinion biases.

2.5. Benchmark Alternatives

The framework is useful only if it differs from simpler alternatives. Three benchmarks discipline the interpretation. The first is a one-index demand model in which a characteristic raises or lowers a single latent preference for RA. Such a model can produce different magnitudes across adoption and conditional share, but it does not naturally predict stable opposite signs without adding selection or boundary effects. The repeated sign reversals on experience, underdiversification, turnover, and confirmation therefore point to a two-cost structure rather than a single RA-taste index.

The second benchmark is generic automation demand. If RA and AUTO were simply different intensities of the same automation preference, then characteristics that predict automation should shift AUTO and RA in the same direction within investor-month. The authority-ladder design tests this directly. A variable that raises the DIRECT-to-AUTO step but not the AUTO-to-RA step is not a generic automation variable; it is associated with execution automation but not full delegation.

The third benchmark is performance chasing. If investors use RA because recent self-directed returns are poor or because RA returns are high, then recent performance should absorb the characteristic patterns. The framework allows returns to matter through Δr_{it}, but it predicts that performance is only one signal among several. The empirical implication is therefore weaker and more realistic: return components can predict RA share, but the main cross-margin signs should remain economically meaningful after performance is examined.

2.6. Authority Transfer Versus Financial Automation

The DIRECT-AUTO-RA product ladder separates automation from delegation. Define

$$c_T = \mathbb{1}\{k \in \text{AUTO}, \text{RA}\}, \qquad c_S = \mathbb{1}\{k = \text{RA}\}, \tag{10}$$

where k indexes the product row in the stacked investor-month design. The coefficient on $X_{it} \times c_T$ identifies how X shifts the investor from DIRECT toward automated execution. The coefficient on $X_{it} \times c_S$ identifies the additional shift from AUTO to full delegation of fund-selection authority. These are product-use differentials within an investor-month rather than literal transition hazards.

This decomposition matters because the automation literature and the delegated management literature speak to different frictions. Algorithm aversion concerns reluctance to rely on an algorithm (Dietvorst et al., 2015, 2018). Delegated management also requires the investor to give up authority over the underlying allocation, a friction closer to the household-finance literature on advice, trust, and delegated portfolio management (Gennaioli et al., 2015; Guiso et al., 2008; Gurun et al., 2018). AUTO and RA allow these margins to be separated within the same investor-month.

Prediction 4 (Authority-step heterogeneity)**.** *A characteristic can have different signs on the automation differential captured by c_T and the full-delegation differential captured by c_S. Skill may be associated with automated execution without increasing full delegation; authority-preserving biases may discourage both differentials; engagement may move investors differently across the two authority layers.*

The prediction is intentionally step-specific. It avoids the stronger claim that all self-directed behavior must reduce RA relative to AUTO. The empirical object is the product-step gradient, not a single generic automation coefficient.

The step-specific formulation is important for interpreting mixed signs. A characteristic may increase the automation differential because it makes scheduled execution attractive, while decreasing or not affecting the full-delegation differential because fund-selection authority remains valuable. Conversely, engagement may be negatively associated with simple automation but positively associated with full delegation if engaged investors prefer products that solve allocation problems rather than

merely automate purchase timing. These patterns would be lost in a binary RA-versus-non-RA comparison.

2.7. Performance Feedback

Performance enters the share equation through Δr_{it}, but the framework distinguishes three return signals. A single level return may be too noisy to move the RA share. Positive and negative return components can instead capture salience or dispersion, because large realized gains and losses both draw attention to the portfolio. Finally, strategy-adjusted RA performance captures allocation within the RA channel, not necessarily demand for the RA channel as a substitute for self-direction.

Prediction 5 (Performance-signal interpretation)**.** *If investors respond to salient return realizations rather than to a single linear return signal, level-return controls can be economically small while positive and negative return components predict conditional RA allocation. Strategy-adjusted RA performance should be interpreted as within-RA selection unless it is separately tied to adoption.*

This prediction limits the performance interpretation. Return evidence can show whether past returns predict allocation, but it does not by itself identify whether the investor is chasing performance, learning about the algorithm, or reacting to portfolio salience.

The performance prediction also clarifies why first-adoption and continuation can differ. Before first RA use, the investor has no individual RA return history. After adoption, both sleeves can generate observed returns, and the same return realization can affect exit, continuation, and conditional share. The lifecycle return tests therefore belong after the baseline demand results, not inside the baseline Heckman.

2.8. Tenure, Experience, and Reclaiming Control

After adoption, two forces can move the RA share in opposite directions. Exposure to the delegated portfolio can raise perceived benefits of RA if the investor learns from the comparison portfolio. Accumulated experience with the self-directed sleeve can instead raise the value of retained authority. Let E_i^{pre} denote fund-investing experience measured before the first RA month. The tenure slope can be written as

$$\frac{\partial s_{it}^*}{\partial \tau} = \rho_A - \rho_C - \rho_E E_i^{pre}, \qquad \rho_A \geq 0, \quad \rho_C \geq 0, \qquad \rho_E > 0. \tag{11}$$

The term ρ_A captures awareness or learning from RA exposure. The term ρ_C captures average reallocation costs or control motives unrelated to pre-adoption experience. The final term captures the incremental authority cost associated with pre-RA fund experience. The sign of the average tenure slope depends on the balance between ρ_A and ρ_C, but the cross-partial with respect to pre-RA experience is negative.

Prediction 6 (Reclaiming control)**.** *If pre-RA fund experience increases the value of self-directed authority, the tenure slope of RA allocation should be lower for high-pre-RA-experience adopters. The average tenure slope is an empirical object that depends on the balance between awareness updating and authority-cost reallocation.*

This is a within-adopter prediction. It does not require all adopters to exit RA. It predicts a change in the allocation share among investor-months in which RA remains active.

The empirical test therefore has two parts. The average tenure coefficient indicates which force dominates for the average active adopter. The interaction between tenure and pre-RA experience tests the authority-cost channel directly. This framing allows the data to show average deepening or average reallocation without changing the model's core mechanism.

2.9. Consequences on Extensive and Intensive Margins

The framework also guides the descriptive consequence evidence. Adoption can affect the self-directed sleeve even if the delegated share is small, because the investor sees a comparison portfolio and may change her own fund choices. The intensive margin can matter separately if a larger RA share changes the remaining self-directed sleeve through substitution, attention, or selection of what remains outside RA.

The current evidence is descriptive rather than causal. Event-time estimates around adoption speak to the extensive margin by showing how self-directed portfolio outcomes move around the transition into RA. RA-share dose-response bins speak to the intensive margin by showing how self-directed-sleeve outcomes vary with the delegated share after RA is already present. The framework predicts that these margins need not move monotonically in the same direction: entry can be associated with discipline in the self-directed sleeve, while very high delegated shares can leave a concentrated residual sleeve outside RA.

This section is deliberately limited. A causal decomposition would require additional assumptions or designs that are not part of the current result set. The framework only licenses a descriptive two-margin reading of the existing event-time and dose-response evidence.

2.10. Selection on Unobservables

The Heckman inverse Mills ratio captures the covariance between unobservables that increase adoption and unobservables that increase the conditional RA share. In this framework, a positive coefficient can arise from latent trust in algorithmic management, platform attachment, or unmeasured affinity for delegated services. The framework does not identify which residual channel is operative. It treats the Mills coefficient as evidence that adoption and delegation share common unobserved determinants beyond measured demographics, skills, and biases.

The residual-selection interpretation is useful because it prevents overloading observed behavioral variables. If the Mills coefficient remains positive after demographics, skills, and biases enter jointly, then observed behavioral heterogeneity is not the whole adoption-intensity link. The framework names trust and platform attachment as plausible residual channels, but it does not claim to identify them directly.

2.11. Economic Magnitudes and Empirical Scaling

The empirical analysis focuses on economic magnitudes, not only signs. Let $\bar{p}$ denote the relevant adoption or product-use rate and let $\bar{s}$ denote the conditional mean delegated share among RA users.

For a covariate X, the economically relevant extensive-margin object is

$$\mathcal{M}_X^D = \frac{\partial \Pr(d=1)/\partial X}{\bar{p}}, \tag{12}$$

and the economically relevant intensive-margin object is

$$\mathcal{M}_X^S = \frac{\partial s^*/\partial X}{\bar{s}}. \tag{13}$$

The signs of these objects classify the mechanism, but their absolute values determine whether the mechanism is economically important. A variable with a statistically significant coefficient that is only one or two percent of the relevant mean is not interpreted in the same way as a variable that moves twenty or fifty percent of the relevant mean.

This scaling discipline applies differently across the paper's designs. In the Heckman tables, selection-stage marginal effects are read relative to the monthly RA-use rate, and outcome-stage coefficients are read relative to the conditional RA-share mean. In the first-adoption table, first-stage marginal effects are read relative to the lower first-adoption rate, so a small-looking percentage-point coefficient can be economically large. In the authority-ladder design, coefficients are product-choice probability differences across the DIRECT-to-AUTO and AUTO-to-RA steps; the economic question is whether the step effect is comparable in size to the main demographic and behavioral slopes. In the bias-trajectory design, the relevant magnitude is cumulative over tenure:

$$\Delta B_i(h) = h\left(\theta_\tau + \theta_H H_i + \theta_G G_i\right), \tag{14}$$

where h is the tenure horizon, H_i marks high pre-adoption bias, and G_i captures engagement. The same monthly coefficient can be economically small for the average adopter and large for high-bias or highly engaged adopters over a one-year horizon.

The framework therefore uses signs to classify candidate mechanisms and relative magnitudes to assess economic importance. This distinction is important for the paper's contribution. The central claim is not that every statistically significant coefficient receives equal weight. The central claim is that economically large coefficients separate across the extensive, intensive, and dynamic margins in ways that a binary-use design would obscure.

2.12. Empirical Implications

The framework yields one core implication and several auxiliary implications. The core implication is the two-margin sign logic: adoption and conditional delegation can have opposite signs because fixed entry costs differ from marginal authority costs. This is the main empirical object of the paper.

The auxiliary implications explain why the paper estimates additional designs. Correctable active-opinion biases should have different signs and post-adoption trajectories from attachment and cognitive-constraint biases. Authority transfer should be distinguished from execution automation using the DIRECT-AUTO-RA product ladder. Pre-RA experience should predict faster reallocation away from RA over tenure if experience raises the value of self-directed authority. Performance tests should separate level returns, positive and negative return components, and within-RA strategy-relative re-

turns. Descriptive consequence evidence should distinguish the extensive margin of adoption from the intensive margin of delegated share. All of these objects should be interpreted relative to their economically relevant base rates or conditional means.

The hierarchy is important. Authority, bias-trajectory, performance, lifecycle, and consequence tests do not compete with the main two-margin result. They discipline its interpretation by asking whether the observed sign patterns are more consistent with decision authority, selective awareness and behavioral correction after exposure, return feedback, lifecycle composition, or residual selection. The framework does not claim unique psychological identification. It specifies which observable patterns would support each mechanism and which patterns would point to an alternative account.

3. Data and Methods

3.1. 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. 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 neighbourhood as Betterment, Wealthfront, and other retail robo-advisors.

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. The third is BangNiTou itself, which we label RA: here the investor hands the fund choice to the algorithm, which selects 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. On the third rung, robo-advising, she delegates the choice of what to own as well. The step that matters for this paper is the one from the second rung to the third. Both rungs automate execution, so only that final step transfers the decision about which funds to hold. Comparing the two therefore isolates one thing: the investor's willingness to cede the decision itself, as opposed to her willingness to automate the mechanical work of trading. This comparison is the basis for the within-investor-month design in Section 4.4.

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 only economic difference between them is who chooses the funds. That is what lets us study the demand for delegating a decision in isolation, rather than confounding it with a general taste for automation.

The setting offers three further advantages. First, the platform records RA assets, self-directed fund holdings, AUTO transactions, cash-management balances, and engagement measures in a single

account. This avoids the external-holding problem that arises when an investor's assets are spread across brokers that the researcher cannot see. Second, all three choices are made inside the same mobile app, so differences between them are unlikely to reflect access frictions or separate distribution channels. Third, the data are monthly and begin in the product's early years, which lets us distinguish first entry, continued use, exit, re-entry, and post-adoption changes in the self-directed portfolio.

This study was remotely conducted on the Ant Open Research Laboratory [1] 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. The main regression variables include basic variables and investment variables.

3.2. 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 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 matter for interpretation, since the paper studies active mobile-platform fund investors rather than 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. 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.

The main estimates are unweighted. Population weights would re-scale the ever-adopter and never-adopter strata back toward their platform shares. In thin monthly cross-sections, though, this mainly amplifies the influence of rare observations and adds little to the slope we are after. The appendix reports weighting checks by stratum, by wealth quintile, and by age cohort, and the headline

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

cross-margin signs for fund experience, underdiversification, total turnover, and confirmation bias survive all of them. We therefore read the main estimates as unweighted cross-sectional slopes, and we do not use them to infer unstratified platform penetration.

Two further features of the design follow from the stratification. First, investors are sampled on ever adoption rather than 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.

3.3. Outcomes and Covariates

Two 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 main extensive-margin outcome is an indicator for holding positive RA assets in month t. The main intensive-margin outcome is

$$s_{it} = \frac{\text{RA Value}_{it}}{\text{RA Value}_{it} + \text{Mutual Fund Value}_{it}}, \qquad d_{it} = 1. \tag{15}$$

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 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.

The lifecycle outcomes separate states that a single binary RA-use measure would pool together. First adoption occurs at a rate of 1.43 percent per investor-month. Continuation, the rate of still holding RA given that the investor held it the prior month, is 94.92 percent. Re-entry among lapsed adopters is 1.25 percent. 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 failed or paused. A single use dummy would blur all three, and the lifecycle analysis therefore treats them separately.

The covariates fall into demographics, platform assets, risk preference, prior robo-product exposure, platform tenure, fund experience, selection ability, timing ability, platform engagement, and nine 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 is a stock of experience, not a measure of recent returns. The distinction matters for our interpretation later on, which turns on accumulated familiarity with selecting and monitoring funds rather than on whether the investor happened to earn high returns. Selection ability and timing ability come from a standard decomposition of realized fund performance. 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, which is why we separate them following Grinblatt and Keloharju (2009).

The nine behavioral biases are central to the analysis and come from separate strands of the behavioural-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, and it serves as the standard proxy for overconfidence, since investors who overestimate the precision of their information trade more than is warranted. A narrower speculative-turnover variant isolates the round trip in which a position is sold soon after a gain and then quickly rebought. Limited attention is the tendency to trade whatever funds are salient, such as those with extreme recent returns or fresh launches, rather than 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 behavioural pattern from its rational counterpart.

Each measure is constructed to capture its own mechanism rather than a single underlying trait. Underdiversification is portfolio HHI, the sum of squared portfolio weights. It equals one for an investor who holds a single fund and falls toward zero as wealth spreads evenly across many funds. We report the effective number of holdings and the weighted return correlation as complementary diversification moments. The measure connects to evidence that retail investors hold concentrated portfolios and often miss the diversification cost (Blume & Friend, 1975; Goetzmann & Kumar, 2008; Huberman, 2001; Polkovnichenko, 2005). Total turnover is the main overtrading proxy, linked to evidence that high-turnover retail investors underperform on average (Barber & Odean, 2000, 2001; Linnainmaa, 2011). 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). It is the proportion of gains realised minus the proportion of losses realised, and it is positive when an investor sells winners more readily than losers. We compute it on a position panel rooted in the investor's $t-1$

information set. 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, capped so that new purchases do not mechanically inflate it. Herding measures trading in funds with high lagged aggregate flow, with a leave-one-out correction in the spirit of Lakonishok et al. (1992).

Two 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 feed directly into the taxonomy in Section 4.6, which depends on telling correctable active-opinion biases, attachment biases, and cognitive-constraint biases apart. A noisier omnibus behavioral index would not support that classification.

Continuous regressors are winsorised at the first and ninety-ninth percentiles and standardised within month. Binary variables are left unstandardised. 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.

3.4. 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:

$$\Pr(d_{it} = 1 \mid Z_{it-1}) = \Phi(Z'_{it-1}\gamma). \tag{16}$$

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

$$s_{it} = X'_{it-1}\beta + \delta\widehat{\lambda}_{it} + \varepsilon_{it}, \qquad d_{it} = 1, \tag{17}$$

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 (**<empty citation>**). Standard errors are Newey–West with lag three. 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 35 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 right way to study a decision whose second margin is realised only after the first. The inverse Mills coefficient is, in turn, an economic object rather than a nuisance term. Its sign tells us whether the investors who select into RA on unobserved dimensions go on to delegate more or less once they enter, and the results section reads it that way.

The Fama–MacBeth implementation is substantive for a related reason. Many of the 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.3 reports a panel fixed-effects control-function alternative, in which the behavioral intensive-margin signs are unchanged while the demographic gradients are, by construction, absorbed into the investor fixed effect. Before each monthly estimation, the pipeline drops variables that are constant within the month, checks rank conditions, removes perfectly separating regressors in the probit, and skips any month whose likelihood fails to converge. These guards bind mainly in the richer bias specifications, where the rolling-window variables reduce the number of usable early months.

Platform tenure is the excluded variable. The exclusion rests on a discovery channel. Longer platform tenure should raise the probability of encountering and evaluating RA, but conditional on adoption and the controls, it should not directly move the share delegated to RA. As in every Heckman design, this assumption is not directly testable, and it is the main identifying restriction behind the selection correction. The relevance condition is strong: months since registration enters the selection equation with coefficients of 0.012, 0.008, and 0.014 in the baseline, skills, and biases panels, with t-statistics of 10.54, 7.80, and 11.05. The exclusion concern is that tenure could instead proxy for platform trust or comfort with digital finance, and so affect the delegated share directly. Two considerations limit this concern. The main conclusions rest on the stability of the observed slopes and cross-margin signs, not on the Mills coefficient alone. And the appendix reports a sensitivity check that adds platform tenure to the outcome equation, where the substantive slopes barely move, which bounds how much any direct-tenure channel could matter.

Selection-stage coefficients in the non-interacted Heckman tables are reported as approximate marginal effects, using

$$\mathrm{ME}_k = \phi\big(\Phi^{-1}(\bar{P})\big)\,\hat{\gamma}_k, \tag{18}$$

where $\bar{P}$ is the relevant monthly event rate. The main Heckman uses panel-specific scaling factors of 0.360 in the baseline panel, 0.368 in the skills panel, and 0.379 in the biases panel. Outcome-stage coefficients are reported in raw RA-share units. A coefficient of 0.020 in the outcome equation therefore means a two percentage-point increase in the conditional RA share, about 13 percent of the conditional mean of 0.151.

3.5. Mechanism and Lifecycle Designs

Four further designs build on the baseline Heckman. The first is the authority-ladder design. It 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. 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. These are within-investor-month product-use differentials, not literal transition hazards.

The second is the reclaiming-control test. It 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. Pre-RA experience is measured over the six months before first RA use and standardised across eligible adopters. Tenure is the cumulative number of RA-active months, demeaned within calendar month. The design asks a single question: do investors who had more fund-market experience before adopting cut their delegated share faster as tenure accumulates?

The third is the bias-trajectory test. It tracks behavioral measures over RA tenure, with rolling six-month trailing means of underdiversification, turnover, and the comparison biases as the headline outcomes. 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, 47,151 adopters out of 66,647 investors in total, rather than on the full 100,000-investor 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 design is descriptive. What it asks is whether the biases behind the selective-awareness interpretation decline after RA exposure.

The fourth set adds return predictors. These designs 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 and examines above-median return indicators. The lifecycle-return table studies first adoption, exit, and continuation separately.

A final set of consequence designs uses event-time and dose-response specifications to separate extensive- and intensive-margin associations. The event-time figures describe changes around adoption. The dose-response table relates lagged RA-share bins to self-directed-sleeve outcomes. Neither is a causal treatment effect. 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.

We probe the main results along two dimensions. The first is the variable side. Alongside the primary delegated-share outcome, we estimate a secondary risky-RA-share outcome that excludes low-risk RA strategies and money market funds from both the numerator and the denominator, which isolates substitution among genuinely risky products. We also re-estimate the intensive margin with two alternative outcomes, the log of RA assets and RA assets over total financial assets, to confirm that the main signs are not an artefact of the RA-plus-mutual-fund denominator. And on the entry side, we reweight the choice-based sample with calibrated stratum, wealth-quintile, and age-cohort

weights, since the headline estimates are unweighted.

The second is the method side, where we vary the specification and the estimation sample. We estimate the Heckman model with demographics, skills, and biases entered together rather than in nested blocks. We estimate a stand-alone AUTO Heckman to confirm that AUTO and RA have distinct determinants. We re-run the authority-ladder design on the full stacked panel and under alternative pre-window lengths. We re-estimate the reclaiming-control test under alternative adopter-sample restrictions and with a fractional-logit link in place of OLS. And we report stability across sample splits and adoption cohorts. Inference throughout uses the Fama–MacBeth time series of monthly coefficients with Newey–West standard errors.

4. Who Uses Robo-Advising? Baseline Evidence and Mechanism Tests

This section first documents the baseline answer to who uses RA. It then lets the observed sign patterns motivate the mechanism tests that follow. The logic is simple. When a characteristic carries opposite signs on adoption and delegated share, the two margins must reflect different fixed and marginal costs, and each later test interrogates one reading of those costs. One test asks whether the relevant marginal cost is decision authority rather than automation. Another asks whether experienced investors reduce their delegated share over tenure. A third asks whether the behavioral sign patterns correspond to selective awareness, attachment, and cognitive constraints. A fourth asks whether recent returns, rather than stable characteristics, predict the two margins.

Throughout, 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. 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, these benchmarks describe the estimation sample, not the platform-population prevalence.

4.1. Baseline Heckman Results

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 nine behavioral measures. The three panels retain 48, 47, and 35 monthly cross-sections.

Wealth is primarily an entry variable. A one-standard-deviation increase in log total financial assets raises RA participation by 14.1 percentage points in the baseline panel, 13.4 in the skills panel, and 8.0 in the biases panel. Measured against the table-specific RA-active rates, these are roughly 43, 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 -0.077 and -0.047, which is about 51 and 31 percent of the 0.151 conditional mean. But once behavioral controls enter, it falls to -0.008, about 5 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 raises both margins. The participation coefficients run from 1.9 to 2.9 percentage points, about 5 to 9 percent of the relevant participation mean. The conditional-share coefficients run from 1.0 to 2.2 percentage points, about 7 to 15 percent of the conditional mean. RA is therefore not simply a low-risk default for cautious investors. More risk-tolerant investors are both more likely to enter and more willing to allocate a larger share once inside. This makes risk preference a useful benchmark for the rest of the table, because it moves both margins in the same direction. It shows what a non-reversing characteristic looks like, which in turn signals that the reversals we find below for experience and the biases are features of those variables rather than a mechanical property of the two-stage design.

Female investors are less likely to use RA and allocate a smaller conditional share, in every panel. The extensive-margin coefficient is about four to five percentage points, roughly 12 to 14 percent of the participation mean, and the conditional-share coefficient is about two percentage points, roughly 13 percent of the conditional mean. Prior use of other robo products on the platform is the strongest observed predictor of RA participation. Its extensive-margin coefficient is about 22 percentage points in every panel, between a half and two-thirds of the participation mean, and its conditional-share coefficient rises to 8.2 percentage points in the biases panel, more than half of the conditional mean. This variable captures substantial persistent heterogeneity in willingness to use delegated products, heterogeneity that measured wealth, risk preference, and 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 -0.006. After behavioral controls enter, age becomes a positive adoption predictor, at 1.7 percentage points per within-month standard deviation, while its conditional-share coefficient is essentially zero. The reason is that age bundles offsetting behavioral channels. Online Appendix Table A.16 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 5 percent of the participation mean in the biases panel. This is a direct explanation for why binary adoption studies reach different conclusions about age: 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 fits a marginal authority cost, since an investor with a large stock of fund-investing experience gives up something valuable

when she cedes fund-selection authority. The joint specification in Online Appendix Table A.1 keeps the positive participation coefficient 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 8 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 6.6 percentage points and raises the conditional RA share by 2.0 percentage points. The entry effect is about 18 percent of the biases-panel participation mean, and the share effect about 13 percent of the conditional mean. The informative part is the sign pair. 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 uniquely predicts: the investors who would benefit most are slower to adopt, yet delegate heavily once a diversified algorithmic portfolio shows them what they had been underestimating. Turnover tells the same story with the weight shifted. It lowers participation by 3.1 percentage points, about 8 percent of the participation mean, and raises the conditional share by 4.1 percentage points, about 27 percent of the conditional mean. This is the same negative-then-positive signature, but tilted toward intensity rather than entry, the reverse of underdiversification. Overtrading deters trying the product less, plausibly because high-turnover investors are already active, but it drives how much they hand over more, since the rule-based product substitutes most directly for the discretionary trading the measure captures. Both signs survive in the joint appendix specification, with coefficients of -0.065 and $+0.016$ for underdiversification and -0.028 and $+0.039$ for turnover. 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.

Confirmation bias has the reverse signature. It raises participation by 2.2 percentage points and lowers the conditional RA share by 1.2 percentage points. These effects are smaller than the underdiversification and turnover effects but still interpretable, about 6 percent of the participation mean and 8 percent of the conditional mean. The joint specification gives $+0.019$ and -0.011. The pattern fits an attachment interpretation, and its positive-then-negative ordering is the exact mirror image of the awareness biases, which is what makes it interpretable. 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. This reading connects the RA demand evidence to the disposition-effect and belief-confirmation literatures (Bailey et al., 2011; Odean, 1998; Shefrin & Statman, 1985).

Attention, gambling, and herding biases reduce both margins. Their magnitudes are smaller than the main active-opinion effects, but the same-signed pattern is itself informative, precisely because it does not reverse. A bias that lowers both margins is one that exposure to the product does not relax. The underlying constraint travels with the investor into the robo sleeve rather than being corrected by it, so there is no margin on which recognition can show up. The split between this same-signed group and the reversing 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). Within this group the details vary. Anchoring has a positive adoption coefficient and no reliable conditional-share coefficient, extrapolation is negative on adoption and small on intensity, and the disposition effect is small throughout. The lesson is that behavioral biases are heterogeneous in their demand signatures: correctable active-opinion biases, attachment-type biases, and cognitive-constraint biases do not respond to RA 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 is a conditional entry variable rather than a stable intensity variable. 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.

4.2. Selection on Unobservables

Selection on unobservables is positive in all main panels. The Mills coefficient is 0.186 in the baseline panel, 0.230 with skills, and 0.251 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.

The positive sign is the economically telling part. It means that the marginal adopter, the investor drawn in by something the covariates do not capture, delegates more than the average adopter rather than less. Two readings can be ruled out on that basis. A cautious-experimenter reading would predict a negative coefficient, since a toe-in-the-water entrant would delegate less than average. A pure adverse-selection reading would also point negative. The positive sign fits neither. What it fits is a single positively valued disposition that operates on both margins at once, most naturally a latent comfort with delegated algorithmic management.

The stability of the coefficient is itself informative. Adding fund experience, selection ability, timing ability, and nine behavioral measures does not absorb the selection channel. 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.

The natural interpretation is latent trust or comfort with delegated algorithmic management. That reading is anchored in the broader trust-in-finance literature, where trust predicts participation beyond wealth, education, and measured sophistication (Guiso et al., 2008), and where delegated management can offer reassurance rather than only alpha (Gennaioli et al., 2015). The platform setting makes durable attachment plausible too, since a user who trusts Alipay as a financial interface may be more willing to try RA and to place a larger share inside it. The design does not separate latent trust, platform attachment, and a general affinity for delegated services. We therefore treat the trust language as an interpretation of a residual covariance, not as a directly identified psychological variable.

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.

4.3. Why the Baseline Requires Mechanism Tests

The baseline table is informative but not decisive. Opposite signs across the two margins are consistent with several stories. They could arise from authority costs, from behavioral self-correction, from a generic demand for automation, from performance feedback, or from lifecycle composition. The remainder of the section therefore treats Table 2 as a diagnostic map rather than a terminal result. Experience and self-directed behavior motivate the authority-ladder and reclaiming-control tests. Underdiversification, turnover, and confirmation motivate the three-channel bias taxonomy and the post-adoption bias-trajectory tests. The possibility that recent returns are the omitted state variable motivates the performance-prediction tests. The baseline answers who uses RA; the mechanism tests ask whether that use reflects automation, authority transfer, behavioral correction, or performance feedback.

These auxiliary analyses do not replace the baseline fact. They discipline its interpretation. A generic-automation reading would predict similar product gradients for AUTO and RA. A pure-performance reading would let recent returns absorb the characteristic gradients. A no-learning reading would predict little post-adoption movement in the biases RA is meant to discipline. Each test puts one such alternative under pressure while preserving the central two-margin object.

4.4. Decision Authority Rather Than Financial Automation

The baseline reading of the intensive margin rested on an authority cost. Investors with more to lose from ceding fund selection, the experienced ones, delegated a smaller share. But that reading assumes the thing investors hesitate over is the transfer of a decision. A simpler alternative is that RA demand just reflects a general taste for 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 3 separates them. The design stacks each investor-month into DIRECT,

AUTO, and RA product rows. The coefficient γ_T captures the product-use differential associated with automated execution relative to direct investing. The coefficient γ_S captures the additional differential associated with RA's full delegation of fund-selection authority. Investor-month fixed effects absorb wealth, market conditions, a general automation affinity, and every other investor-month characteristic common across the product rows. The language of steps is therefore an authority-ladder interpretation of within-investor-month product differentials, not a literal sequential transition model.

The stand-alone AUTO Heckman in Online Appendix Table A.4 explains why this comparison is necessary. AUTO has its own selection and intensity relations. Wealth and fund experience raise AUTO participation but lower AUTO intensity, and the AUTO Mills coefficient is negative rather than positive. The behavioral contrasts are sharper still. Underdiversification raises AUTO intensity by 8.8 percentage points, more than four times the RA conditional-share coefficient in Table 2, while turnover is negative on both AUTO margins. The two products cannot be treated as interchangeable forms of automation. They attract related but distinct behavioral margins, which is what makes the within-investor product ladder the right test of decision authority.

The results are hard to reconcile with a simple generic-automation reading. In the composite specification, the skill index raises the automation step by 2.3 percentage points but has no significant effect on the full-delegation step. Figure 2 plots the two step coefficients with their confidence intervals and the equality tests, and the visual fact is that the covariates do not shift the DIRECT-to-AUTO and AUTO-to-RA steps in parallel.

Skill, then, moves investors from direct investing toward automated execution, not toward the additional RA delegation. Authority-bias lowers both differentials, by -0.041 on the automation component and -0.015 on the full-delegation component. Comparison bias raises automation by 2.4 percentage points but has no full-delegation effect. Platform engagement carries opposite signs across the two, lowering automation by 4.5 percentage points while raising full delegation by 1.7 percentage points. These splits are on the order of the demographic effects in the baseline Heckman. The engagement reversal is the telling one, because a generic-automation account cannot produce it. A common taste for automation would push both steps the same way. That engaged investors decline mere scheduling but embrace full delegation says the two steps solve different problems. Scheduling automates a chore the active investor is content to keep; full delegation hands off the allocation decision she may have concluded she cannot win. The two product differentials are therefore economically distinct choices, not two labels on a single automation preference.

The economic content is concrete. A skilled investor can value automated execution because it reduces transaction hassle while still preserving her right to choose the funds, yet need not value RA's additional transfer of fund-selection authority. An engaged investor can do the reverse, disliking 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 rather than about digital product use in general.

The disaggregated rows reinforce the point without requiring a uniform self-direction story for every variable. Selection and timing ability load on the automation step and are close to zero on full delegation. Turnover-related variables generally discourage product-step movement, while some

comparison biases matter only for one step. The feature that matters is not that every covariate shares a sign. It is that the signs and magnitudes vary across the automation and full-delegation steps in a way a single use indicator cannot show.

The contribution is that RA demand is not a generic preference for digital automation. The distinctive margin is the transfer of fund-selection authority. The result links the household-finance literature on delegated management and trust to the algorithm-aversion literature on willingness to rely on algorithms (Dietvorst et al., 2015, 2018). The field setting adds a margin that laboratory experiments typically cannot observe, since the same investor can choose direct fund-picking, automated execution with authority preserved, or full delegation.

This clarifies what the paper adds to both literatures. The delegated-advice literature often studies whether investors trust an adviser or intermediary enough to participate. The algorithm-aversion literature often studies whether people prefer human or algorithmic forecasts after observing errors. The Alipay setting puts a related but different choice in the data, namely whether the investor keeps or transfers fund-selection authority when execution automation is already available. The point is not only that some investors avoid algorithms, but that the decision right being transferred can be isolated from automation itself.

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.

4.5. Reclaiming Control Over Tenure

The authority-cost reading carries a dynamic implication that the baseline cross-section cannot test. 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 re-learns 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 test therefore 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. Table 4 does this. The outcome is the RA share among RA-active investor-months. The key regressor is pre-RA fund experience, measured over the six months before first RA use, interacted with RA tenure. If experience raises the value of self-directed authority, then experienced adopters should reduce their RA share more quickly over tenure.

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. Online Appendix Table A.9 preserves the sign 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 and grows more negative as pre-adoption experience rises. That gradient is what makes the interaction matter economically.

The economic size is large enough to bear 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 non-trivial 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.

The result complements the authority ladder. The ladder isolates 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 arc appears from the self-directed side in the event-study evidence of Section 6.1, where portfolio concentration falls sharply at adoption and then recovers toward its pre-adoption level over the following months. 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.

The tenure test also clarifies the role of experience in the main Heckman. Experience is not only 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 lifecycle pattern is closer to the learning-by-trading interpretation in which fund-market participation builds an outside option: the investor can delegate, but she continues to value the authority to manage the self-directed sleeve.

4.6. A Three-Channel Taxonomy of Biases

It helps to separate the levels of the argument before reading the signs. At the level of the 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 of these. It resolves into three sign signatures, which are sub-mechanisms within the behavioral channel rather than independent paper-level claims. The first sub-mechanism is awareness. Underdiversification and turnover are

mistakes RA can plausibly discipline, but investors with more severe mistakes may not perceive the cost before adopting, so they enter less yet delegate more once exposed. 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 sub-mechanisms is what distinguishes the behavioral evidence from a uniform account in which all biases respond to RA in the same way.

The baseline left the awareness reading underdetermined. 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. If exposure to a diversified, rule-based portfolio is what makes the cost of these mistakes salient, then the mistakes themselves 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. Table 5 tests exactly this. 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 the Online Appendix 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, 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. These are large changes relative to the scale of the underlying measures, especially for the high-bias groups. For turnover this tenure decline also completes the arc left open at first adoption, where turnover deterred entry but did not yet raise the initial delegated share (Section 5.1). 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, and the high-bias and active-opinion-taking interactions remain negative and significant. Engagement is the most useful auxiliary discriminator here. The reason is that the engagement interaction, rather than the level decline, is what 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. The Adopter-by-tenure-by-engagement triple interaction is negative and significant for both focal biases in the smoothed panel, and remains negative and significant in the raw monthly panel. For the focal pair, the raw-monthly engagement coefficients are -0.0003 for underdiversification and -0.0007 for turnover. The comparison biases do not show the same systematic engagement pattern, which is the right contrast, since a general post-adoption drift would dose every bias and does not. Some comparison interactions are statistically significant, 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.

Figure 4 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.

The result connects the demand-side paper 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.

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. The table is informative for a comparative reason. 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.

4.7. Performance as a Predictor of the Two Margins

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 coefficients might be proxies for performance feedback. This subsection is therefore an alternative-explanation check, not a fourth paper-level mechanism alongside selective awareness, decision authority, and reclaiming control. The question is narrow: can recent returns explain away the two-margin characteristic patterns? If they could, the paper would be less about fixed costs, authority,

and behavioral self-correction, and more about return feedback. The evidence below leaves room for performance salience while making the stronger performance-only reading hard to sustain.

Appendix Tables A.5 and A.6 show that single level measures of fund and RA performance are economically small in the RA-share equation. Fund-performance coefficients are no larger than 0.003 in absolute value. Category-adjusted fund returns at the three- and six-month horizons are statistically significant, but the economic scale stays small. RA excess and beta-adjusted returns are also small and mostly insignificant. Strategy-adjusted RA performance is positive, with coefficients between 0.007 and 0.008. This is best read as within-RA strategy selection: investors allocate more to RA strategies that outperform comparable RA strategies, not necessarily to RA because it beats the market.

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 together they are 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, which says the same thing 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 flips. 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. The near-symmetry is informative, because the rival readings each predict 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 and significant at all three horizons, 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 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 rather than sign on both margins. The level effects, a few percentage points on adoption and well under one percentage point

on the conditional share, are modest next to the wealth, experience, and bias coefficients, and they leave the larger opposite-sign patterns on those characteristics intact.

The practical implication is that performance belongs in the paper as an alternative channel and a boundary condition. 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. But in both places the effect is modest, and it does not overturn the main interpretation that stable investor characteristics and behavioral states separate adoption from conditional delegation.

The result is also not a contradiction of 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.

5. First Adoption and Lifecycle Stages

The pooled monthly Heckman combines first entry, continuation, re-entry, and other active states into a single participation measure. That pooling is the reason it cannot say what first brings an investor to RA, as opposed to what characterizes her once she is an active user. This section separates first adoption from the other lifecycle states, and then asks whether performance predicts the different margins in different ways.

5.1. First Adoption of RA and AUTO

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 8 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 rather than her contemporaneous active-user state.

Keeping the stock and the flow straight is what turns apparent sign reversals 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: 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. One clarification about that 7.7 percent figure is needed here, because it is higher than the 1.43 percent monthly first-adoption transition rate in Table 1. The two numbers measure different things. 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 rather than 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. This mirrors the pooled Heckman: 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 and has no reliable initial-share effect. Female investors show a two-margin pattern in miniature. 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.

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. This is the sharpest case of the stock-versus-flow distinction, because it contrasts with the positive pooled RA participation coefficient in Table 2, and the contrast 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 and raises the authority cost of ceding fund selection, so 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.

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 suggestive rather than decisive. 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. A toe-in-the-water entrant who then deepens delegation as comfort accumulates would generate exactly this reversal 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 4.5 rather than 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 the survivors rather than a recognition response. First adoption shuts that door. 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 cleanest form of the pattern the paper documents. Confirmation reproduces its reverse signature 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 is the instructive partial case. It lowers first RA adoption but has no significant initial RA-share effect, so its entry deterrence replicates while its positive-share reversal does not appear at first use. This is not an inconsistency with the pooled result. It is a statement about timing. Recognition of concentration is immediate, and it is visible in the initial delegated share. The turnover reversal 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 4.6. 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 reinforces the authority interpretation. 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 but not RA initial intensity. Confirmation raises AUTO first adoption and has a positive but insignificant AUTO-share coefficient, whereas it lowers the initial RA share. Execution automation and full delegation are therefore already distinct at the entry point, not only among active users.

5.2. Performance Before First Adoption and Other 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.11 and A.12 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. 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 9 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 4.7 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 rather than 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. 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, rather than 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 mirror image. 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. 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 rather than levels.

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.

6. Consequences on the Extensive and Intensive Margins

The paper's main estimates are demand estimates, not causal treatment effects. Even so, the two-margin framework makes a prediction about consequences that is worth checking. Adoption may change behavior even when the delegated share is small, because the investor now sees an algorithmic comparison portfolio and works through a delegated-management interface. Conditional delegation may have a separate gradient, because a larger RA share changes the residual self-directed sleeve. If entry and commitment are genuinely different margins, their consequences for the self-directed portfolio should look different too. The current result set provides descriptive evidence on this through event-time figures, adopter-panel sleeve regressions, and lagged RA-share dose-response regressions. It does not contain a matched first-time-adopter versus non-adopter treatment design, nor a causal decomposition of total effects into entry and intensity components, so we state the evidence as margin-specific associations.

The test these designs jointly apply is simple. If entry and delegated share were empirically interchangeable, then event-time changes and RA-share dose-response patterns would move monotonically in the same direction. They do not, and the direction in which they fail to is what the rest of the section reads.

6.1. Event-Time Evidence Around Adoption

The event-study figures speak to the extensive margin: 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 shows an HHI decline of about 3.8 percentage points at spell start and a trough around 5.7 percentage points several months after adoption. RA share jumps at adoption and then declines over the first year. Turnover rises in the adoption month, consistent with mechanical portfolio reallocation at entry, and then falls below its pre-adoption reference in the following months. 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 1 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. The effect then attenuates, reaching 0.358 by month 12, 0.381 by month 18, and 0.403 by month 24. 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.

Online Appendix Table A.18 gives a complementary cross-sectional view within the RA-adopter sample. The regressor is the lagged combined-advisor indicator, the platform's separate portfolio-

advisory product rather than the BangNiTou robo-advisor itself. Its coefficient therefore describes how an adopter's self-directed sleeve looks in months when she also holds the combined-advisor product. In those months, adopters hold 8.93 more self-directed funds, have lower self-directed HHI by 0.065, carry a higher self-directed share by 0.022, and trade less in both total and speculative turnover. The association is consistent with platform advisory products coexisting with a broader and more disciplined self-directed sleeve. Its role is descriptive, since delegation does not mechanically displace self-direction.

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.

6.2. Dose Response in Delegation Share

The event study held the size of delegation aside; the dose-response design brings it back. Table 10 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 $(0, 0.20]$, 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 calmer 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.

6.3. 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 RA use can coexist with substantial self-directed investing rather than 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.

7. Synthesis

The evidence supports a unified but deliberately restrained interpretation. The central empirical fact is that robo-advisor demand separates into adoption and conditional delegation. The mechanism evidence then explains why that separation is economically meaningful. Fixed entry frictions govern discovery and trial. Marginal authority costs govern how much wealth is ceded after entry. Behavioral self-correction and attachment determine which biases reverse across the margins. Performance salience and post-adoption reallocation shape allocation inside active states without driving the core result.

The baseline Heckman delivers the two-margin fact. Wealth and age operate mainly on entry once behavioral controls are present, while fund experience and several biases move adoption and conditional delegation in different directions. These patterns extend the household-finance distinction between participation and portfolio efficiency (Calvet et al., 2007; Campbell, 2006) to a delegated product. Here the intensive margin is not a portfolio-efficiency measure but the share of wealth ceded to an algorithm. The canonical alignment between participation and better portfolios need not hold when participation itself requires an authority transfer.

Within that fact, the authority ladder isolates the relevant intensive-margin object. RA is not merely automation. It is the additional step from automated execution to delegated fund selection. Skill raises the automation step but not full delegation, authority-bias reduces both, and engagement moves differently across the two. These results connect the paper to the algorithm-aversion literature while adding a field margin that is hard to observe experimentally. The relevant question is not only whether investors trust an algorithm, but whether they will transfer the specific authority over fund selection.

Authority transfer is dynamic, not static. Among active adopters, longer RA tenure is associated with lower delegated shares, and the negative tenure-share gradient is steeper for investors with more pre-RA fund experience. RA users are not all moving monotonically deeper into delegation. Experienced users can treat RA as a partial or transitional device. The authority cost that appears in the within-month product ladder reappears over time, as investors reallocate back toward self-direction.

On the behavioral side, correctable active-opinion biases reduce adoption, raise conditional delegation, and decline after RA exposure, while attachment and cognitive-constraint biases carry different signatures. The first-adoption evidence is what makes the awareness reading hard to dismiss as a composition artifact, since the positive delegated-share response to underdiversification is present at the

investor's first delegation decision, before survivorship in the active pool can manufacture it. The two focal awareness biases then correct at different speeds. Underdiversification is immediate, raising the initial delegated share at entry. Turnover is gradual, its reversal surfacing among active users and deepening with tenure. This links demand for RA to the corrective-effects literature. The question is not only whether RA improves portfolios conditional on use (Bianchi & Brière, 2026; Hao et al., 2022; Reher & Sokolinski, 2024; Rossi & Utkus, 2024), but which investors choose to expose themselves to the correction. Demand-side selection is therefore potentially welfare-relevant, and relevant for the aggregate gains from corrective technology, even when the product benefits its adopters. We do not quantify those gains.

Performance narrows the leading alternative. Recent returns matter in decomposed specifications and within RA strategies, but single performance controls are too small to explain the main patterns. The evidence fits salience, dispersion, or within-channel selection better than a story in which investors enter RA after poor self-directed returns or raise their RA share only after RA outperformance. It also remains compatible with the algorithm-aversion experiments, because the panel return signal is a noisy sleeve-level outcome rather than a clean, experimentally assigned algorithm-error signal. The consequence evidence sits alongside these demand estimates as supporting rather than identifying material. Entry is associated with changes in self-directed behavior, and higher delegated shares alter the composition of the remaining self-directed sleeve in a non-monotone way.

The interpretation remains observational. The framework labels mechanisms but does not uniquely identify psychological primitives. Selective awareness, attachment, cognitive constraints, authority costs, and latent trust are disciplined by sign patterns and auxiliary tests, not by randomized variation in beliefs or product design. The contribution is the cross-margin regularity, together with the linked set of field tests that make the regularity interpretable.

8. Robustness

The appendix collects robustness checks, organized around the claims they defend. We summarize them here by the part of the argument each one protects.

The first claim is that the cross-margin signs are not an artifact of how the specifications are nested. The joint Heckman in Online Appendix Table A.1 enters demographics, skills, and biases together. It leaves the cross-margin signs for underdiversification, turnover, and confirmation intact. Fund experience continues to predict participation, while its conditional-share coefficient attenuates, which is why the main text reads that intensive-margin relation as partly mediated by behavioral correlates. A second check addresses the estimator itself. Online Appendix Table A.3 reports a panel control-function selection design with investor and month fixed effects and two-way clustered standard errors. It recovers the same intensive-margin behavioral signs, which confirms that the cross-margin patterns do not depend on the Fama–MacBeth estimator or on Newey–West inference.

The second claim is that the headline signs are not an artifact of the RA-plus-mutual-fund denominator. Online Appendix Table A.2 re-estimates the intensive margin with two alternative outcomes, log RA assets and RA over total financial assets. The slope patterns persist. The inverse Mills coefficient, however, turns negative for log RA assets and stays positive for the two share outcomes. That

flip turns on what each outcome measures. The share outcomes ask what fraction of fund wealth the investor cedes, which is a compositional decision. Log RA assets asks how many yuan she places, which is a scale decision that moves with wealth. The latent disposition raises the ceded fraction but not the absolute amount, which says the unobserved factor is a willingness to hand over control rather than a wealth proxy in disguise. Selection on unobservables therefore operates on the compositional delegation-share margin, not on the absolute scale of delegated assets.

The third claim is that AUTO and RA are distinct products rather than two intensities of one. The stand-alone AUTO Heckman in Online Appendix Table A.4 confirms it. Underdiversification carries a much larger positive coefficient on AUTO intensity than on RA intensity, confirmation is positive on both AUTO margins, turnover is negative on both, and the AUTO Mills coefficient is negative in the biases specification. These contrasts support the authority-ladder claim that AUTO is not simply a low-intensity form of RA.

The fourth claim is that the mechanism results are stable under alternative samples and windows. Online Appendix Tables A.7, A.8, and A.9 vary sample restrictions and pre-window lengths for the authority and reclaiming-control tests. The authority patterns hold in the full stacked panel and under alternative pre-window definitions. The reclaiming-control result remains negative under alternative adopter-sample restrictions and in the fractional-logit specification.

The fifth claim is that performance does not absorb the baseline patterns. Online Appendix Tables A.5 and A.6 add the single-performance controls and show that fund-side and broad RA-side level returns are too small to do so, while strategy-adjusted RA returns predict allocation within the RA channel. Online Appendix Table A.13 extends the lifecycle-return design.

Finally, Online Appendix Table A.15 and Figure A.3 report stability across sample splits and cohorts, Online Appendix Table A.16 documents the age-bias correlations used to interpret the age coefficient, and Online Appendix Table A.18 reports self-directed-sleeve composition within the adopter-month panel. These checks support the sign patterns without removing the observational nature of the mechanism tests, which we read as disciplined support for the framework rather than as unique psychological identification.

9. Conclusion

Robo-advising involves the delegation of investment decisions, not merely digital access to financial markets, and this distinction changes how demand should be measured. Adoption and delegation conditional on adoption respond to different costs, and the same investor characteristic can drive the two margins in opposite directions. Therefore, a binary adoption regression obscures the economic object that matters most for delegated portfolio management.

Using Alipay panel data from September 2020 through September 2024, we showed that wealth and age mainly affect entry once variable that control for behavioral patterns are included. Furthermore, fund experience and several behavioral variables exhibit opposite effects on adoption and delegated share, while positive selection on unobservables remains even after extensive controls. Underdiversification and turnover reduce adoption but raise conditional delegation; confirmation bias display the opposite pattern, while attention, gambling, and herding reduce both margins. TThese

patterns help explain the conflicting findings in the robo-advisor adoption literature, where different studies report different demographic or behavioral predictors because they observe different mixtures of entry and commitment.

The mechanism evidence explains why the distinction matters. RA differs from automatic investment plans because it transfers fund-selection authority, and skill increases automated execution without increasing full delegation. Among active adopters, the decline in delegated share over tenure is steeper for investors with more pre-RA fund experience. Behavioral mistakes that robo-advice directly addresses decline after adoption, particularly among initially high-bias and more engaged investors. While performance feedback affects some dimensions of robo-advisor use, it does not account for the main cross-margin patterns associated with investor characteristics. Descriptive evidence on post-adoption portfolio adjustments further shows that adoption and delegation are linked to different changes in self-directed investing. Albeit descriptive, this evidence reinforces the paper's core interpretation.

The policy implication is not simply that robo-advising should be expanded. The demand-side selection problem matters, because the investors who would benefit from correction may not be the most likely to enter, and those who enter may later reallocate back toward self-direction. Product design should therefore distinguish cues that help investors recognize own-portfolio mistakes from generic marketing of automation or recent performance. A platform that wants to reach underdiversified or high-turnover investors may need to make those mistakes salient before adoption; advertising aggregate algorithmic performance is unlikely to substitute for that self-diagnostic channel.

The broader implication is that participation in delegated financial products should be studied as an authority-allocation problem. Target-date funds, managed accounts, advisor-sold funds, and other delegated products all require investors to cede some decision rights. For such products, entry, delegation, and post-adoption reallocation are distinct economic decisions. The central finding of the paper is that these decisions have systematically different determinants once adoption and conditional delegation are modeled separately.

Several limitations remain. The setting is one platform in one country, and the same integrated product menu that enables identification may limit external validity. The key insight, however, is not the level of adoption among Chinese super-app users, which is specific to this market, but the cross-margin logic that underlies delegated investing. A delegated product can involve both a fixed cost of entry and a marginal cost of ceding decision authority, and these costs may move in opposite directions. That structure should apply wherever participation requires delegation. The mechanism tests are observational and provide supporting rather than definitive evidence on the roles of selective awareness, trust, and attention. Likewise, the evidence on how investors adjust their portfolios after adopting robo-advice is descriptive and is not intended to estimate the causal effects of adoption. Future experiments that vary portfolio-composition salience, delegated authority, or trust-related messaging could identify these channels more directly. Nevertheless, the central conclusion remains: demand for delegated portfolio management cannot be understood through adoption alone, but requires separate consideration of adoption, delegation intensity, and post-adoption reallocation.

References

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. (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.

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.

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

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.

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

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

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.

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

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.

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.

Gurun, U. G., Stoffman, N., & Yonker, S. E. (2018). Trust busting: The effect of fraud on investor behavior. *Review of Financial Studies*, *31*(4), 1341–1376.

Hao, R., Hu, C., Xu, X., & Zhang, Y. (2022). Beyond performance: The financial education role of robo-advising. *Available at SSRN 4230191*.

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

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.

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

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

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.

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.

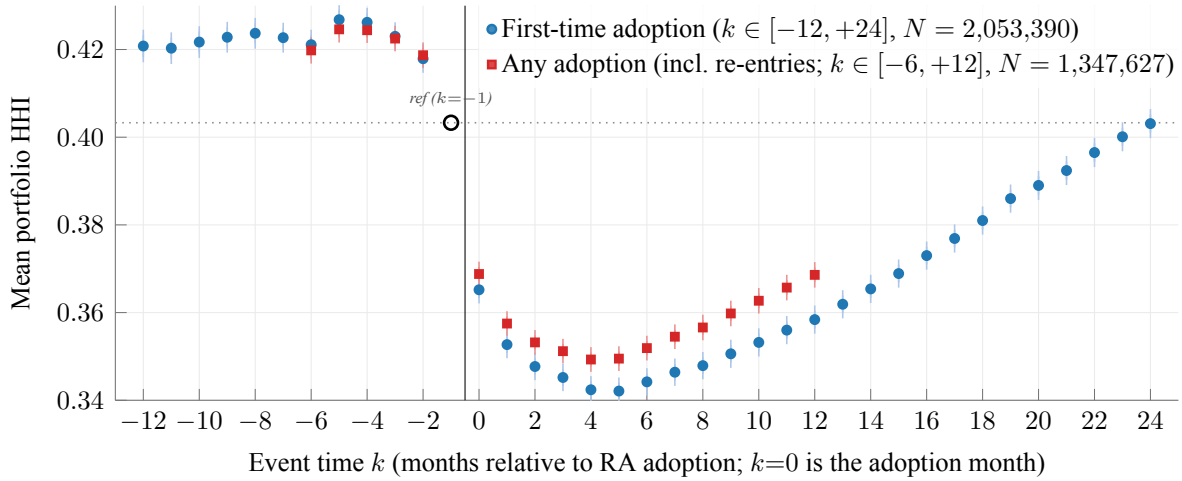
Rossi, A. G., & Utkus, S. P. (2024). The diversification and welfare effects of robo-advising. *Journal of Financial Economics*, *157*, 103869.

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.

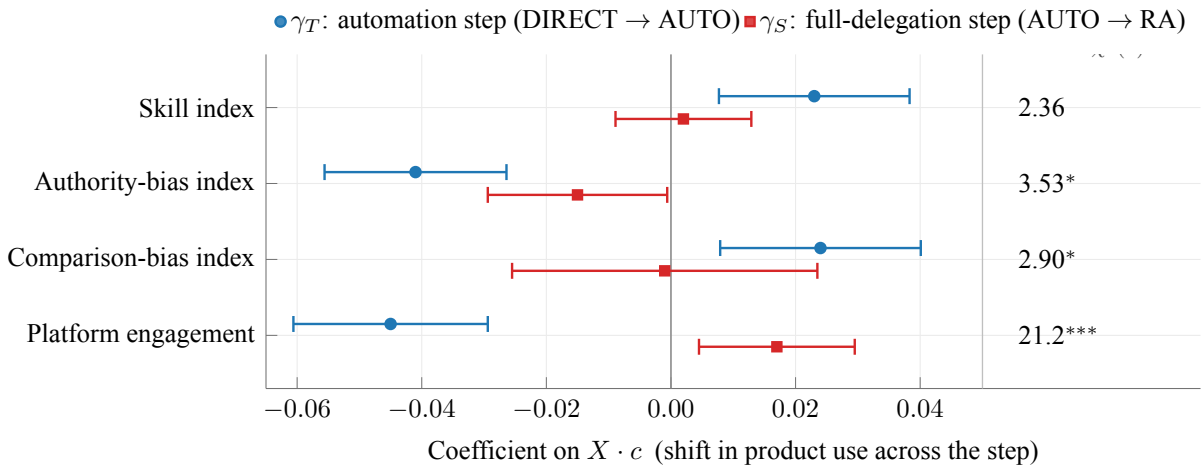
Figures

Figure 1: What Changes When Investors Adopt RA?



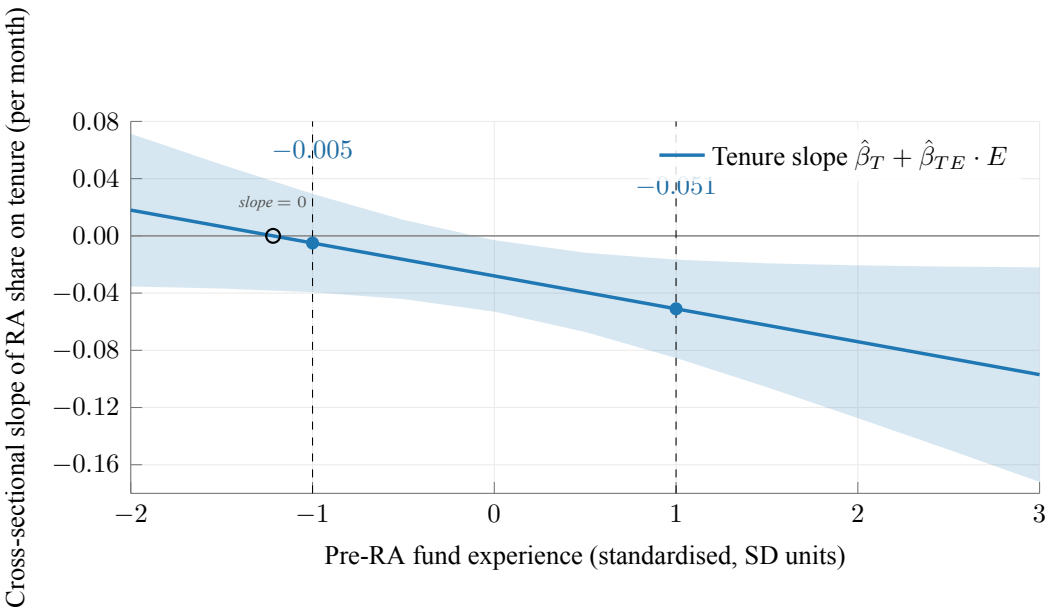
Notes. Mean portfolio Herfindahl index (HHI) by event time relative to RA adoption. Each level is computed as $\bar{y}_{k=-1} + \hat{\beta}_k$, where $\hat{\beta}_k$ is the event-study coefficient from a within-investor fixed-effects regression of HHI on event-time dummies and $\bar{y}_{k=-1}$ is the 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 2: Decision-Authority Mechanism: What Moves Investors Across Each Delegation Step



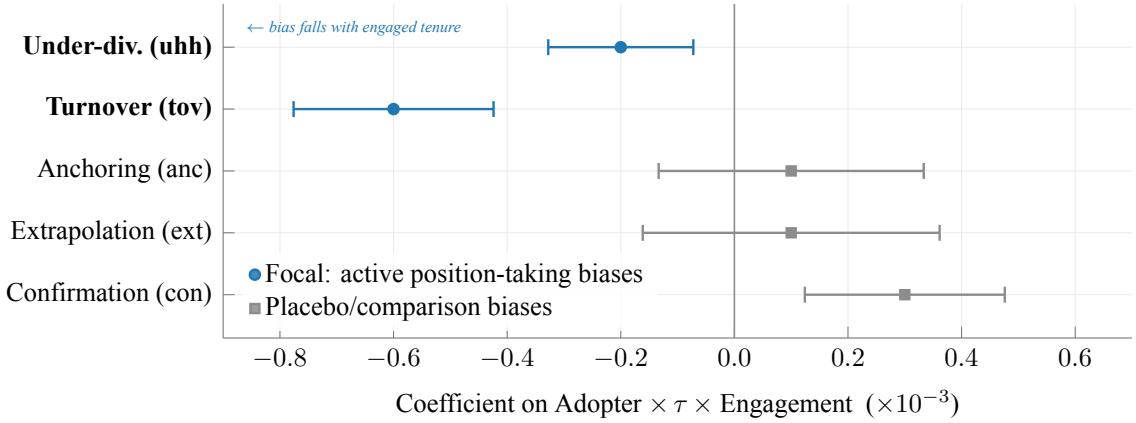
Notes. Each row reports two coefficients from the headline stacked linear-probability specification in Specification (1) of Table 3, estimated on the investor-month × product panel with investor-month fixed effects. γ_T (blue circles) is the coefficient on $X \cdot c_T$ and identifies how covariate X moves the investor across the *automation step* (DIRECT → AUTO, where the platform automates purchase timing); γ_S (red squares) is the coefficient on $X \cdot c_S$ and identifies the additional shift across the *full-delegation step* (AUTO → RA, where the algorithm also selects securities). Horizontal bars are 95% confidence intervals; standard errors are recovered as $|\text{coef}/t|$ from t-statistics two-way clustered on investor and month. The right-hand column reports the Wald $\chi^2(1)$ test of $\gamma_T = \gamma_S$ for each index ($^{*}\ p < 0.10$, $^{***}\ p < 0.01$); overlap of the two confidence bars is not the equality test, which uses the coefficients' joint covariance. *Reading the figure.* Skill loads on the automation step only ($\gamma_T = +0.023$) and is flat at the full-delegation step ($\gamma_S = +0.002$): more skilled investors are associated with automated timing but not with additional fund-selection delegation. The authority-bias index is negative at both steps, most strongly at automation. Platform engagement shows the sharpest split—strongly negative at the automation step ($\gamma_T = -0.045$) yet positive at the full-delegation step ($\gamma_S = +0.017$), the only index whose two steps differ at the 1% level (Wald $\chi^2 = 21.2$). The comparison-bias index loads on automation only, consistent with a non-specific intensity channel rather than a delegation-specific one.

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



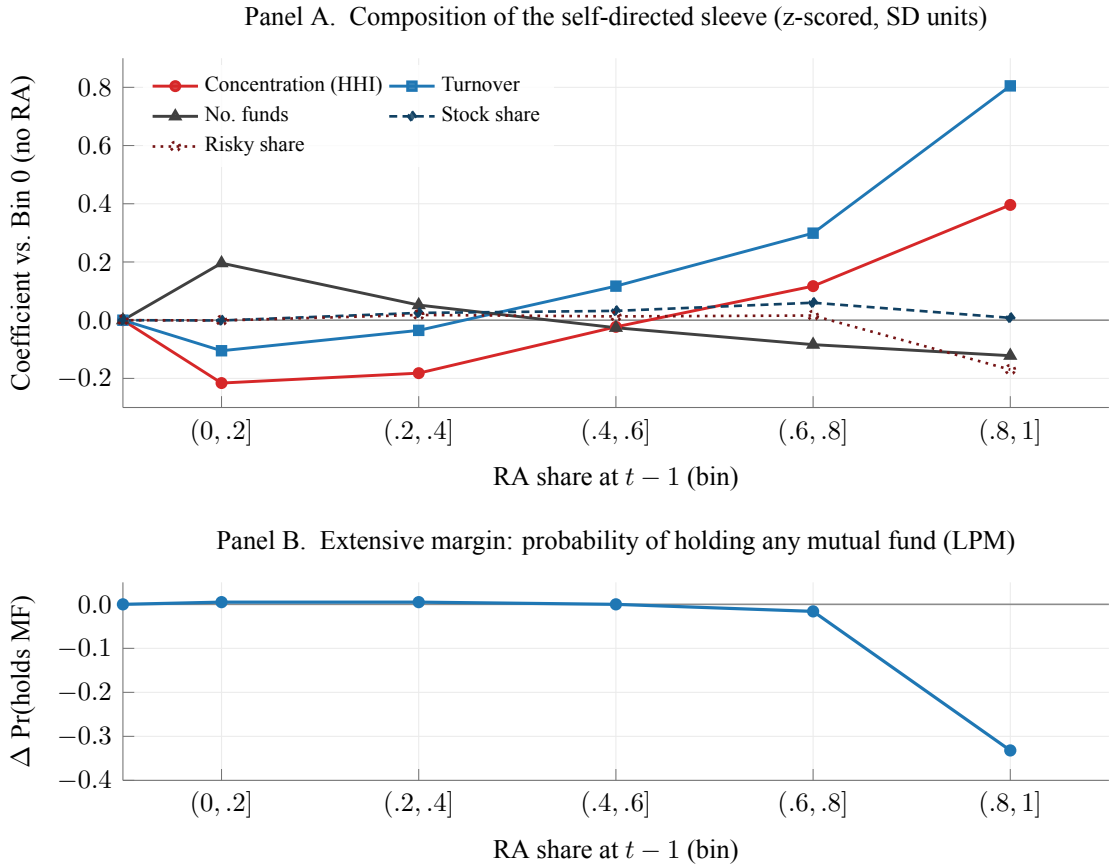
Notes. The figure plots the cross-sectional slope of RA share (RA assets / (RA + MF), among RA-active adopter-months) on robo-advisor tenure, as a function of standardised pre-RA fund experience E. Because estimation is Fama–MacBeth (monthly cross-sectional regressions), this is the average within-month 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 4 (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: 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 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 10, 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
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%; raw levels, 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%; raw levels)								
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
Panel E. Auto-Invest Usage								
Auto-Invest Usage (monthly)	4,900,000	0.323	—	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
Panel F. Adoption Lifecycle (Investor-Month Transitions)								
Month-to-month state transitions across 4,900,000 investor-months:								
Stay out (non-RA → non-RA)	3,248,698	66.30%	—	—	—	—	—	—
Continue (RA → RA)	1,491,648	30.44%	—	—	—	—	—	—
Enter (non-RA → RA)	79,827	1.63%	—	—	—	—	—	—
Exit (RA → non-RA)	79,827	1.63%	—	—	—	—	—	—

Continued on next page

Table 1 (continued)

Variable	N	Mean	SD	Min	p25	Median	p75	Max
Conditional rates:								
First-adoption rate (per investor-month)	—	1.43%	—	—	—	—	—	—
Continuation rate (| RA at $t-1$)	1,571,475	94.92%	—	—	—	—	—	—
Re-entry rate (| RA at $\leq t-2$ only)	783,622	1.25%	—	—	—	—	—	—

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 statistics are winsorised at the 1st and 99th percentiles but reported in their raw, unstandardized levels.** The within-month standardization is applied separately inside the regressions, not here.

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. Login Count (last 12 months) is reported here for the engagement specifications; the median investor logs in roughly $\sim$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 12-month rolling market β), and (iv) category-adjusted (self-directed funds only) or strategy-peer-adjusted (RA only). **Panel E** reports auto-invest usage: the monthly investor-level participation rate (0.323) and, conditional on participation, the share of fund purchases executed via automated plans (0.733). **Panel F** reports adoption-lifecycle transitions across consecutive investor-months. Of the 4,900,000 panel rows, the four mutually exclusive month-to-month transitions sum to the full panel; rates below convert these counts into conditional rates (continuation conditioned on holding RA last month; re-entry conditioned on having held RA in some earlier month but not last month). 70.0% of investors adopt RA at some point during the sample window (70,034 of 100,000).

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=35)	
Variable	Probit ME	Outcome (RA_{share})	Probit ME	Outcome (RA_{share})	Probit ME	Outcome (RA_{share})
Constant	-0.197***	0.023	-0.190***	-0.065	-0.150***	-0.068
	(-6.745)	(0.381)	(-7.655)	(-1.226)	(-12.626)	(-1.419)
Log(Total Financial Assets)$_{t-1}$	0.141***	-0.077***	0.134***	-0.047***	0.080***	-0.008
	(12.559)	(-4.369)	(11.754)	(-4.735)	(6.036)	(-1.182)
Risk Preference	0.029***	0.010**	0.022***	0.022***	0.019***	0.014***
	(5.616)	(2.493)	(4.261)	(5.806)	(6.626)	(9.505)
Age	0.004	-0.006**	0.003	-0.010*	0.017***	-0.001
	(0.770)	(-2.069)	(0.612)	(-1.678)	(25.518)	(-0.438)
Female	-0.042***	-0.016***	-0.043***	-0.023***	-0.052***	-0.018***
	(-17.144)	(-3.944)	(-19.983)	(-5.083)	(-19.617)	(-4.186)
Combined RA Dummy$_{t-1}$	0.224***	0.045***	0.222***	0.061***	0.219***	0.082***
	(24.914)	(2.955)	(29.157)	(4.590)	(98.397)	(4.989)
Baby Assets$_{t-1}$	-0.037***	0.013***	-0.034***	0.008**	-0.020***	-0.001
	(-18.857)	(3.298)	(-16.497)	(2.493)	(-9.378)	(-0.185)
Fund Experience (Months)$_{t-1}$			0.035***	-0.028***		
			(9.546)	(-5.977)		
Selection Ability$_{t-1}$			0.004*	-0.000		
			(1.886)	(-0.067)		
Timing Ability$_{t-1}$			0.001	-0.001		
			(0.968)	(-1.335)		
Anchoring Bias$_{t-1}$					0.022***	0.002
					(12.976)	(1.244)
Total Turnover$_{t-1}$					-0.031***	0.041***
					(-14.565)	(6.245)
Attention Bias$_{t-1}$					-0.009***	-0.006***
					(-6.870)	(-4.519)
Disposition Effect$_{t-1}$					0.001*	-0.001**
					(1.751)	(-2.009)
Underdiversification HHI$_{t-1}$					-0.066***	0.020***
					(-99.536)	(3.435)
Extrapolation Bias$_{t-1}$					-0.005***	-0.002
					(-2.731)	(-1.467)
Gambling Bias$_{t-1}$					-0.009***	-0.004***
					(-6.546)	(-3.167)
Confirmation Bias$_{t-1}$					0.022***	-0.012***
					(29.161)	(-3.447)
Herding Bias$_{t-1}$					-0.006***	-0.004***
					(-7.322)	(-3.894)
Months Since Registration$_{t-1}$	0.012***		0.008***		0.014***	
	(10.538)		(7.800)		(11.047)	
Inverse Mills Ratio (λ)		0.186***		0.230***		0.251***
		(4.425)		(5.398)		(5.108)
Province Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2 / R^2	0.082	0.132	0.078	0.146	0.090	0.222
N	4,799,616	1,564,155	4,471,044	1,539,119	3,214,908	1,199,373

Notes. Estimates are from unweighted Heckman two-stage (Tobit-2) selection models 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 $\bar{P}_s$ is the average monthly RA-active rate in each specification's estimation sample (RA-adopter months ÷ selection months, on the 4.9-million investor-month panel) and $\hat{\beta}_k$ is the Fama–MacBeth average Probit coefficient: $\bar{P} = 0.326$ (Baseline), 0.344 (+Skills), 0.373 (+Biases), with scaling factors $\phi(\Phi^{-1}(\bar{P}_s)) = 0.360, 0.368$, and 0.379 respectively. Outcome-stage coefficients are reported as estimated (no scaling). Newey–West HAC t-statistics (max lag = 3) in parentheses; 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). All right-hand-side variables except Age, Female, and Risk Preference are lagged one period. Province dummies are included but not reported. 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 3: Decision-Authority Mechanism: Stacked Linear Probability Model with Entity and Time 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 platform engagement			
Skill index	0.023***	0.002	2.364
	(2.946)	(0.360)	[0.124]
Authority-bias index	-0.041***	-0.015**	3.526*
	(-5.522)	(-2.044)	[0.060]
Placebo-bias index	0.024***	-0.001	2.901*
	(2.920)	(-0.080)	[0.089]
Platform engagement	-0.045***	0.017***	21.229***
	(-5.658)	(2.657)	[0.000]
Specification (2). Disaggregated authority-relevant skills and biases			
Selection ability	0.013*	0.003	0.532
	(1.657)	(0.417)	[0.466]
Market timing	0.020**	-0.001	2.245
	(2.324)	(-0.134)	[0.134]
Total turnover	-0.065***	0.003	14.120***
	(-5.573)	(0.414)	[0.000]
Underdiversification HHI	0.039***	-0.008	7.916***
	(3.404)	(-0.966)	[0.005]
Disposition effect	0.002	-0.006	0.397
	(0.316)	(-0.898)	[0.528]
Gambling bias	-0.019*	-0.008	0.419
	(-1.873)	(-1.576)	[0.518]
Speculative turnover	-0.028***	-0.013**	1.615
	(-5.069)	(-1.982)	[0.204]
Specification (3). Disaggregated placebo biases			
Anchoring bias	0.018**	0.025***	0.267
	(2.310)	(3.590)	[0.606]
Attention bias	-0.031***	0.022***	14.568***
	(-4.039)	(3.051)	[0.000]
Extrapolation bias	0.015*	-0.019***	5.388**
	(1.839)	(-2.581)	[0.020]
Confirmation bias	0.068***	0.003	21.524***
	(8.034)	(0.382)	[0.000]
Herding bias	0.010	-0.019***	4.848**
	(1.293)	(-2.973)	[0.028]
N	598,029	598,029	598,029
R^2	0.135	0.137	0.137

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 $k \in \{\text{DIRECT, AUTO, RA}\}$ in month t and zero otherwise; every investor-month is stacked into three rows, one per product alternative.

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 (the *automation step*), and γ_S on $X \cdot c_S$ identifies the additional shift across the AUTO-to-RA step (the *full-delegation step*). The Wald column tests $H_0 : \gamma_T = \gamma_S$ from the clustered covariance matrix.

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; full estimates are reported in the Online Appendix. Estimation is by `linearmodels.PanelOLS` 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. The 598,029 stacked observations cover 199,343 unique investor-months from 10,951 investors across 33 calendar months.

Cluster-robust t-statistics appear in $(\cdot)$ below coefficient estimates; Wald p-values appear in $[\cdot]$ below the corresponding $\chi^2(1)$ statistic. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed). Stars on the Wald column refer to the p-value of the $\chi^2(1)$ statistic.

Table 4: 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 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	$+0.0004$
	(-3.49)	(-11.94)	$(+1.96)$	(-0.40)	$(+0.53)$

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)
$\tau\times$ 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: 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 9: 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).

Table 10: 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).

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 cross-margin logic. Delegated financial products can have fixed entry costs and marginal authority costs, and the same investor characteristic can move the adoption margin and the conditional-delegation margin in opposite directions.

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 main baseline and skills specifications therefore use more monthly cross-sections than the bias-augmented and joint specifications. All continuous regressors are winsorised 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. 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 depends on the covariate vector. 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, rather than the platform-level prevalence of robo-advisor use. Population-calibration weights would re-scale 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. Under random sampling within strata, the slope coefficients remain informative for the cross-sectional relationships of interest, while the intercept is not a target parameter.

We nevertheless re-estimate the main Heckman specifications under alternative weighting schemes. The first uses stratum inverse-probability weights. The second calibrates by wealth quintile. The third calibrates by ten-year age cohort. The headline sign patterns are unchanged: fund experience remains a positive participation predictor with a weaker or negative conditional-share relation; underdiversification and turnover preserve the negative-adoption, positive-delegation signature; and confirmation preserves the reverse signature. The weighted estimates therefore support the conclusion that the main 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 rather than 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,

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

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, $1/\mathrm{HHI}_{it}$, 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 skill and timing skill decompose realized mutual-fund performance into two pieces. Selection skill is the value-weighted excess return of the investor's held funds relative to Morningstar category benchmarks. Timing skill 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.1 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 and positive on conditional delegation; turnover remains negative on adoption and positive on conditional delegation; and confirmation remains positive on adoption and negative on conditional delegation. Fund experience retains a positive participation coefficient but its conditional-share coefficient attenuates 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, which is why the paper treats selection on unobservables as an economic object rather than as a nuisance term fully absorbed by the observed controls.

Online Appendix Table A.2 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. The inverse Mills coefficient changes across outcomes, however. It is positive on the main compositional share, smaller on the total-asset share, and negative on the log-level outcome. This pattern is informative: the positive selection channel is strongest when delegation is measured as a composition of the investor's active fund wallet rather than 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.3 re-estimates the conditional-share equation with a completely different design: a Wooldridge (1995) and Semykina–Wooldridge (2010) control-function panel selection model. 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 is informative about whether the signs are an artifact of either choice.

The 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, 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 Fama–MacBeth and Newey–West machinery of the main specification.

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

Online Appendix Table A.4 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, confirmation is positive on AUTO entry and not negative on AUTO intensity, turnover is negative on both AUTO margins, 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 Tables A.7 and A.8 report authority-ladder robustness checks. The full stacked panel uses all investor-months rather than only the identifying-only sample, while the pre-window-sensitivity table varies the number of months used to construct pre-RA experience and pre-RA behavioral measures. The distinction between the automation step and the full-delegation step is preserved. Skill and engagement do not move the two steps in the same way, which supports the main text's claim that execution automation and delegation of fund-selection authority are different revealed choices.

Online Appendix Table A.9 reports reclaiming-control robustness under alternative adopter-sample restrictions. The baseline specification studies RA-active adopter-months and relates delegated share to RA tenure, pre-RA fund experience, and their interaction. The robustness table changes the adopter sample and estimates fractional-logit variants. Across these variants, the pre-RA-experience by tenure interaction remains negative. 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.9 reports additional bias-trajectory specifications over robo-advisor tenure. 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 rather than as unique identification of awareness.

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 sub-mechanisms 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.5 adds single fund-performance measures to the RA Heckman. The outcome-stage coefficients are economically tiny, 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.

Online Appendix Table A.6 adds RA-performance measures. Broad RA excess returns and beta-adjusted returns are small, while strategy-adjusted RA returns are positive. 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.11, A.12, and A.13 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.

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

Online Appendix Table A.13 complements the main dose-response figure by reporting additional self-directed mutual-fund sleeve outcomes and comparison checks. The outcome variables describe the residual self-directed sleeve after lagged RA exposure: number of mutual funds held, portfolio concentration, turnover, speculative turnover, and related composition measures. 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 pattern is intentionally described as descriptive. Lower and moderate RA shares are associated with broader and calmer 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 the fund wallet 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.15 and Figure A.3 report stability checks across early and late sample windows and adoption cohorts. These exercises are scope checks rather than substitutes for the main 49-month Fama–MacBeth estimates. 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 main sign patterns are evaluated in the longer monthly panels and in the joint specification, not in six-month subsamples alone.

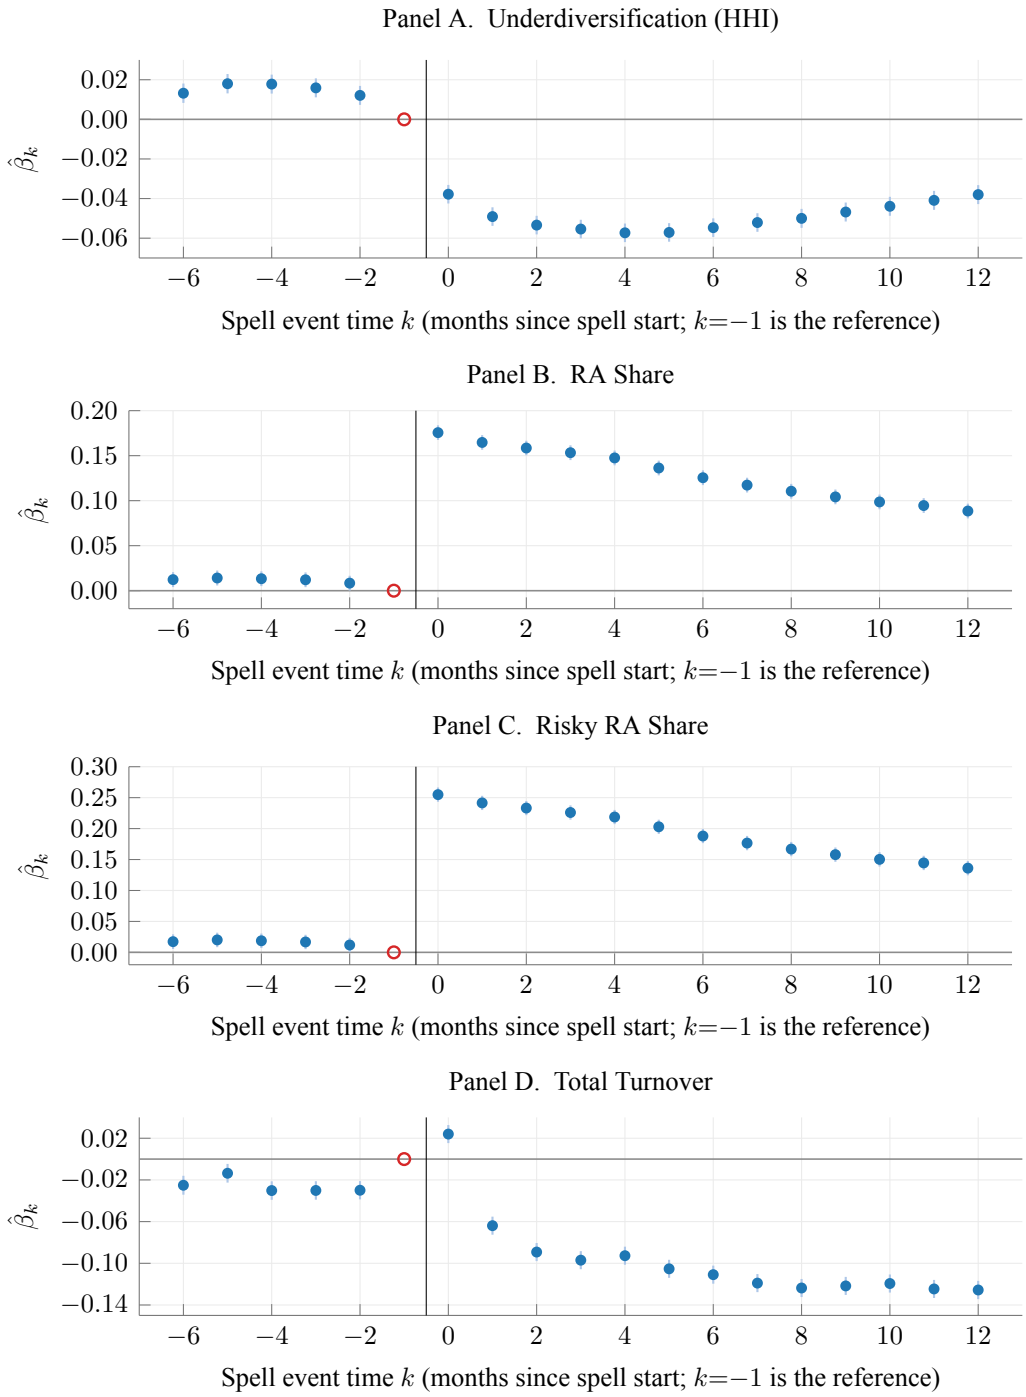
Online Appendix Table A.16 reports correlations between age and the behavioral-bias, skill, and experience measures. These correlations support the interpretation of the age coefficient in the main text. Older investors trade less, are less concentrated by HHI, and display stronger confirmation bias. 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

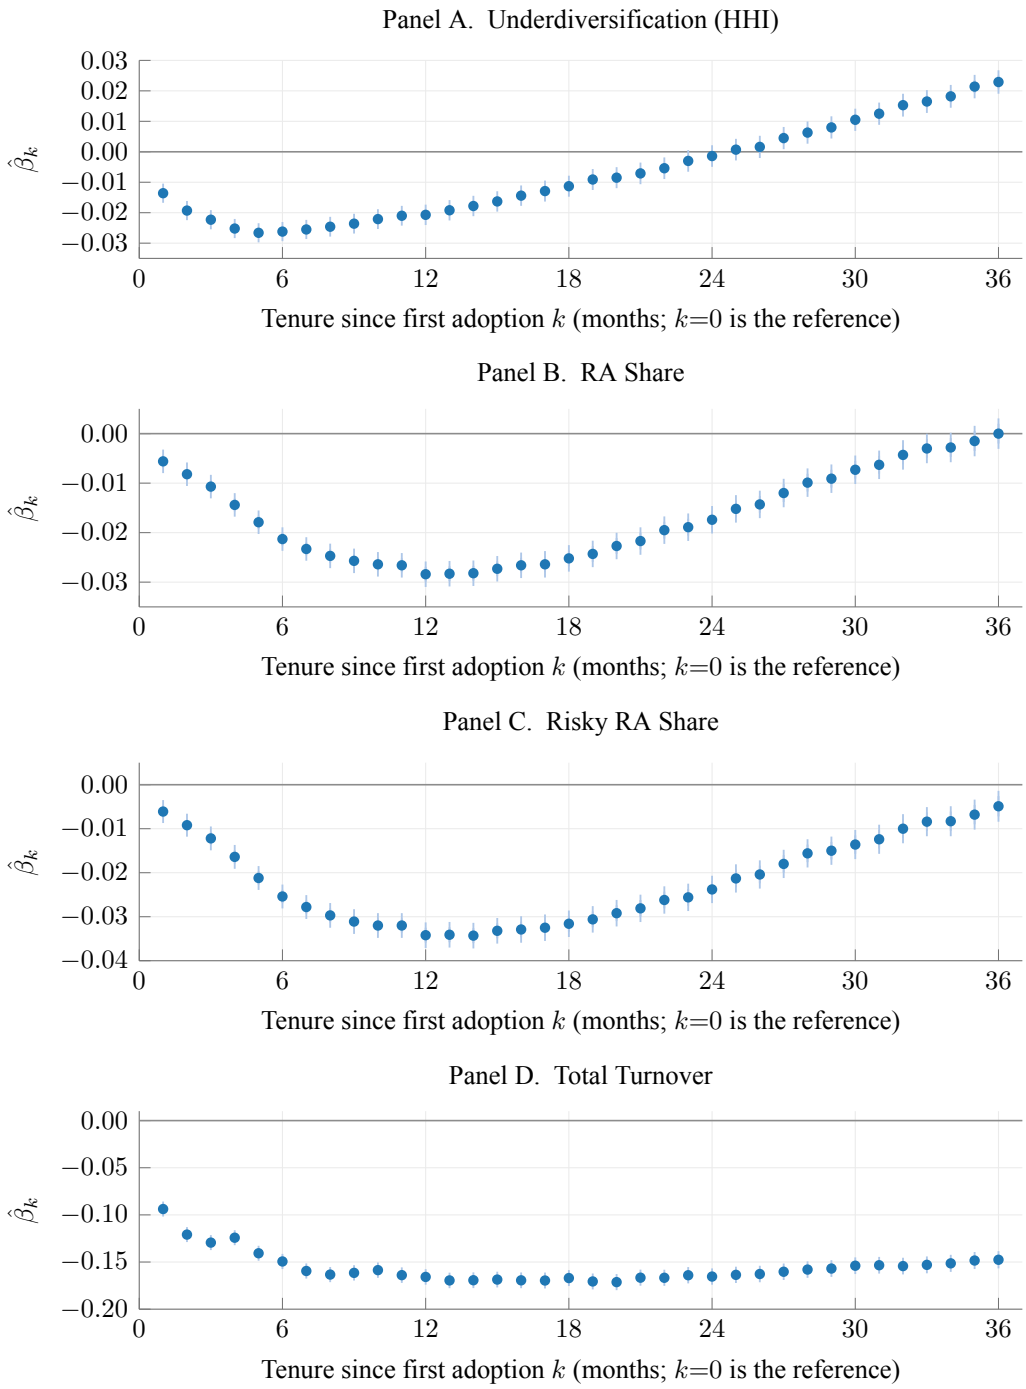
Figure A.1 repeats the event-study analysis for all adoption spells, including re-entry spells, rather than only first-ever adoption. Figure A.2 reports the post-adoption tenure event study for allocation and self-directed-sleeve outcomes. Figure A.3 plots time-stability and cohort-split coefficients. 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)



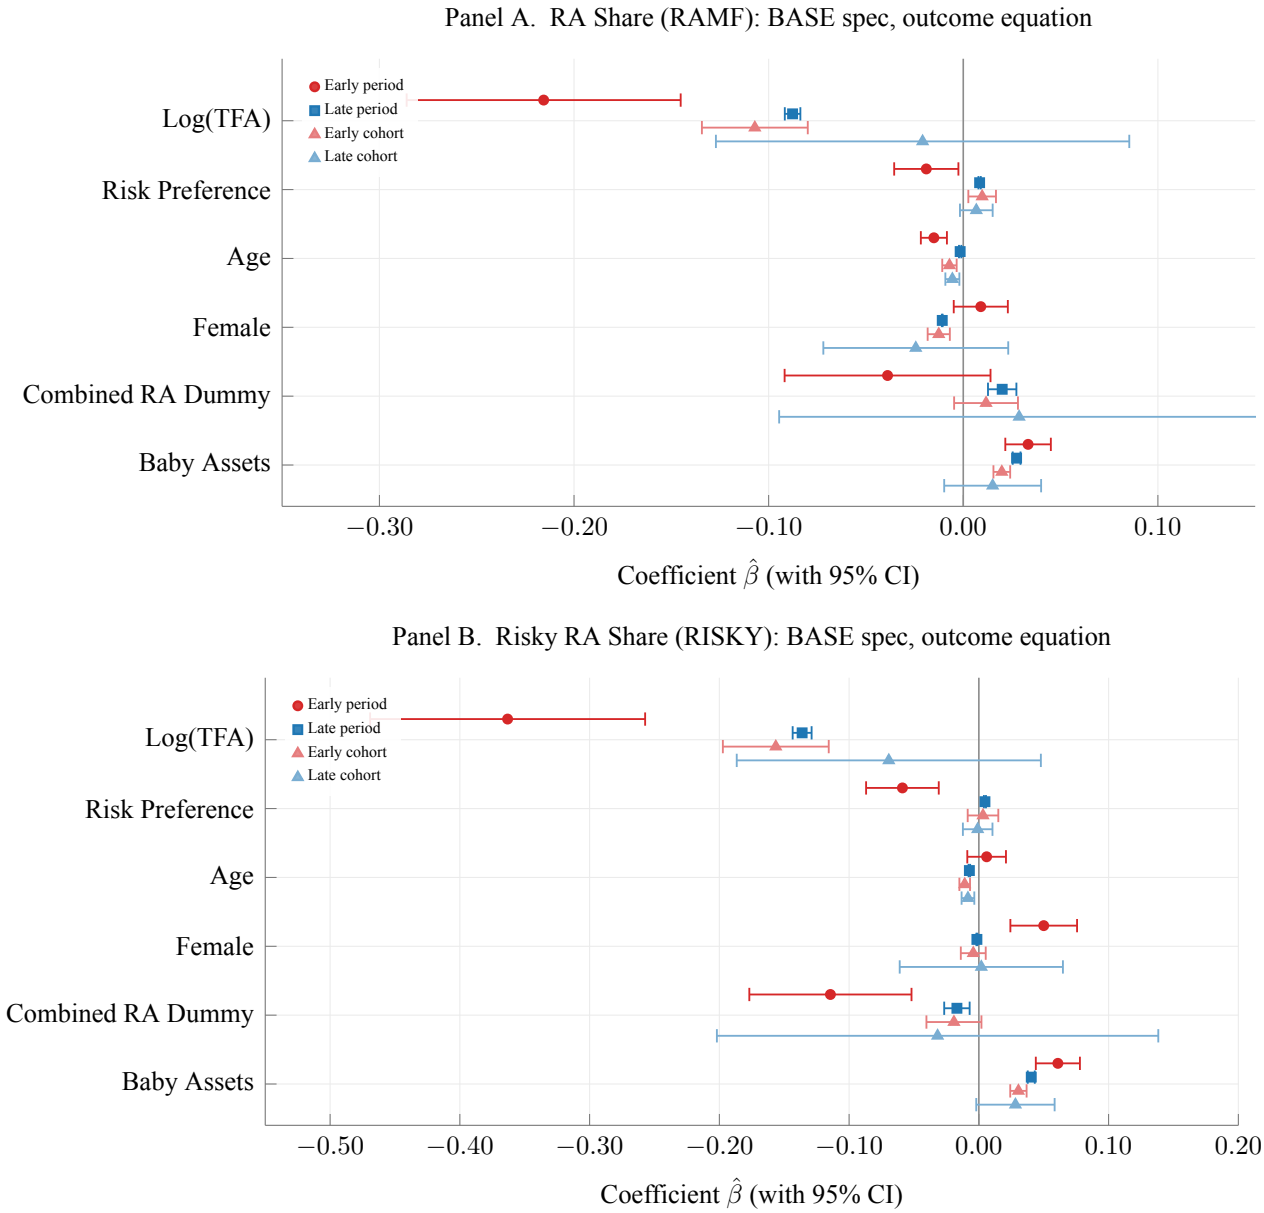
Notes. Event-study coefficients from **spell** fixed-effects regressions $y_{it} - \bar{y}_s = \sum_k \beta_k \cdot \mathbf{1}\{k_{st} = k\} + \varepsilon_{it}$, where s indexes RA-adoption *spells* (every $0 \to 1$ transition of RA_{it} opens a new spell). Demeaning is done within spell rather than within investor, so investors who exit and re-enter RA contribute multiple observations. Window is shorter than the first-adoption event study in Figure 1 ($k \in [-6, +12]$ 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 1) 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 1, 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: 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.2: 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.3: 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.3 (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.4: 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 3 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 3 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.5: 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.6: 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.5 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.5), 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.7: Decision-Authority Mechanism on the Full Panel: Sample-Restriction Robustness

Variable	Automation step DIRECT → AUTO (γ_T)	Full-delegation step AUTO → RA (γ_S)	Equality test $\chi^2(1)$
Specification (1). Composite indices and platform engagement			
Skill index	0.020***	0.003	2.167
	(2.859)	(0.421)	[0.141]
Authority-bias index	-0.036***	-0.016**	2.633
	(-5.202)	(-2.285)	[0.105]
Placebo-bias index	0.021***	0.000	2.291
	(2.670)	(0.021)	[0.130]
Platform engagement	-0.043***	0.017***	22.373***
	(-5.687)	(2.903)	[0.000]
Specification (2). Disaggregated authority-relevant skills and biases			
Selection ability	0.011	0.003	0.400
	(1.566)	(0.524)	[0.527]
Market timing	0.018**	-0.001	2.086
	(2.261)	(-0.100)	[0.149]
Total turnover	-0.058***	0.003	12.957***
	(-5.325)	(0.354)	[0.000]
Underdiversification HHI	0.038***	-0.010	9.478***
	(3.529)	(-1.431)	[0.002]
Disposition effect	0.002	-0.005	0.399
	(0.298)	(-0.929)	[0.528]
Gambling bias	-0.019**	-0.008	0.498
	(-1.993)	(-1.377)	[0.480]
Speculative turnover	-0.025***	-0.011*	1.829
	(-5.199)	(-1.867)	[0.176]
Specification (3). Disaggregated placebo biases			
Anchoring bias	0.016**	0.023***	0.319
	(2.212)	(3.618)	[0.572]
Attention bias	-0.029***	0.021***	14.328***
	(-4.005)	(3.040)	[0.000]
Extrapolation bias	0.013*	-0.017**	4.780**
	(1.701)	(-2.476)	[0.029]
Confirmation bias	0.062***	0.002	20.697***
	(7.667)	(0.270)	[0.000]
Herding bias	0.008	-0.016***	3.750*
	(1.082)	(-2.690)	[0.053]

Notes. This table reproduces the headline specification of Table 3 on the full panel rather than the identifying-only sample. The full panel retains investor-months in which the investor is on the same product across all three stacked rows; the headline drops them. All other modelling choices match Table 3. The dependent variable, products, indicators c_T and c_S, definitions of the automation step and the full-delegation step, the construction of baseline controls, the entity fixed effect, and the clustering structure are as in Table 3. The sample contains 215,729 unique investor-months stacked into 647,187 observations, drawn from 10,956 investors and 33 calendar months. Within-R^2 equals 0.124, 0.125, and 0.126 for Specifications (1), (2), and (3). Cluster-robust t-statistics appear in $(\cdot)$ below coefficient estimates; Wald p-values appear in $[\cdot]$ below the corresponding $\chi^2(1)$ statistic. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed).

Table A.8: Sensitivity to Pre-Window Length L (Identifying-Only Sample)

Variable	$L = 3$	$L = 6$	$L = 9$	$L = 12$
Specification (1). Composite indices and platform engagement				
Skill index (γ_T)	0.016**	0.023***	0.013*	0.014
	(2.151)	(2.946)	(1.678)	(1.634)
Skill index (γ_S)	-0.000	0.002	0.005	0.000
	(-0.053)	(0.360)	(0.720)	(0.043)
Skill index, Wald $\chi^2(1)$	1.671	2.364	0.369	0.890
	[0.196]	[0.124]	[0.544]	[0.345]
Authority-bias (γ_T)	-0.043***	-0.041***	-0.043***	-0.043***
	(-6.121)	(-5.522)	(-5.599)	(-5.267)
Authority-bias (γ_S)	-0.010	-0.015**	-0.012	-0.006
	(-1.484)	(-2.044)	(-1.474)	(-0.730)
Authority-bias, Wald $\chi^2(1)$	6.914***	3.526*	4.798**	6.099**
	[0.009]	[0.060]	[0.029]	[0.014]
Placebo-bias (γ_T)	0.019**	0.024***	0.023***	0.021**
	(2.523)	(2.920)	(2.815)	(2.479)
Placebo-bias (γ_S)	0.006	-0.001	-0.007	-0.008
	(0.919)	(-0.080)	(-0.939)	(-0.980)
Placebo-bias, Wald $\chi^2(1)$	1.015	2.901*	4.214**	3.437*
	[0.314]	[0.089]	[0.040]	[0.064]
Platform engagement (γ_T)	-0.047***	-0.045***	-0.046***	-0.036***
	(-6.158)	(-5.658)	(-5.608)	(-4.508)
Platform engagement (γ_S)	0.015**	0.017***	0.017***	0.013*
	(2.572)	(2.657)	(2.596)	(1.837)
Engagement, Wald $\chi^2(1)$	23.772***	21.229***	20.868***	12.244***
	[0.000]	[0.000]	[0.000]	[0.000]
Specification (2). Disaggregated authority-relevant skills and biases				
Selection ability (γ_T)	0.011	0.013*	0.010	0.004
	(1.525)	(1.657)	(1.249)	(0.443)
Selection ability (γ_S)	0.001	0.003	0.001	-0.001
	(0.111)	(0.417)	(0.091)	(-0.094)
Selection, Wald $\chi^2(1)$	0.666	0.532	0.442	0.089
	[0.415]	[0.466]	[0.506]	[0.765]
Market timing (γ_T)	0.011	0.020**	0.008	0.016
	(1.400)	(2.324)	(0.947)	(1.586)
Market timing (γ_S)	-0.001	-0.001	0.005	-0.001
	(-0.170)	(-0.134)	(0.767)	(-0.186)
Timing, Wald $\chi^2(1)$	0.857	2.245	0.042	1.110
	[0.355]	[0.134]	[0.838]	[0.292]
Total turnover (γ_T)	-0.060***	-0.065***	-0.064***	-0.062***
	(-5.473)	(-5.573)	(-5.149)	(-4.892)
Total turnover (γ_S)	0.005	0.003	0.006	0.008
	(0.613)	(0.414)	(0.686)	(0.858)
Turnover, Wald $\chi^2(1)$	14.028***	14.120***	12.799***	11.277***
	[0.000]	[0.000]	[0.000]	[0.001]
Underdiversification (γ_T)	0.036***	0.039***	0.038***	0.031**
	(3.464)	(3.404)	(3.085)	(2.521)
Underdiversification (γ_S)	-0.007	-0.008	-0.008	-0.001
	(-0.937)	(-0.966)	(-0.972)	(-0.079)
Underdiv., Wald $\chi^2(1)$	8.134***	7.916***	6.552**	2.740*
	[0.004]	[0.005]	[0.011]	[0.098]
Disposition effect (γ_T)	-0.002	0.002	0.000	0.000
	(-0.235)	(0.316)	(0.041)	(0.002)
Disposition effect (γ_S)	-0.005	-0.006	-0.004	-0.007
	(-0.761)	(-0.898)	(-0.667)	(-0.929)
Disposition, Wald $\chi^2(1)$	0.056	0.397	0.127	0.213
	[0.814]	[0.528]	[0.722]	[0.645]
Gambling bias (γ_T)	-0.026***	-0.019*	-0.023**	-0.022**
	(-2.758)	(-1.873)	(-2.090)	(-2.025)
Gambling bias (γ_S)	-0.002	-0.008	-0.004	-0.001
	(-0.327)	(-1.576)	(-0.737)	(-0.128)

Continued on next page

Table A.8 (continued)

Variable	$L = 3$	$L = 6$	$L = 9$	$L = 12$
Gambling, Wald $\chi^2(1)$	2.292	0.419	1.074	1.409
	[0.130]	[0.518]	[0.300]	[0.235]
Speculative turnover (γ_T)	-0.025***	-0.028***	-0.027***	-0.024***
	(-4.655)	(-5.069)	(-4.505)	(-3.662)
Speculative turnover (γ_S)	-0.013**	-0.013**	-0.011*	-0.011
	(-2.184)	(-1.982)	(-1.647)	(-1.508)
Spec. turnover, Wald $\chi^2(1)$	1.135	1.615	1.561	0.947
	[0.287]	[0.204]	[0.212]	[0.331]
Specification (3). Disaggregated placebo biases				
Anchoring bias (γ_T)	0.017**	0.018**	0.016**	0.014*
	(2.363)	(2.310)	(2.017)	(1.669)
Anchoring bias (γ_S)	0.025***	0.025***	0.025***	0.027***
	(3.843)	(3.590)	(3.354)	(3.161)
Anchoring, Wald $\chi^2(1)$	0.363	0.267	0.375	0.576
	[0.547]	[0.606]	[0.540]	[0.448]
Attention bias (γ_T)	-0.032***	-0.031***	-0.030***	-0.027***
	(-4.124)	(-4.039)	(-3.762)	(-3.087)
Attention bias (γ_S)	0.020***	0.022***	0.018***	0.011
	(2.649)	(3.051)	(2.593)	(1.488)
Attention, Wald $\chi^2(1)$	13.053***	14.568***	11.590***	6.008**
	[0.000]	[0.000]	[0.001]	[0.014]
Extrapolation (γ_T)	0.012	0.015*	0.010	0.004
	(1.598)	(1.839)	(1.210)	(0.489)
Extrapolation (γ_S)	-0.016**	-0.019***	-0.018**	-0.017**
	(-2.186)	(-2.581)	(-2.393)	(-2.095)
Extrapolation, Wald $\chi^2(1)$	3.976**	5.388**	3.520*	1.809
	[0.046]	[0.020]	[0.061]	[0.179]
Confirmation bias (γ_T)	0.062***	0.068***	0.067***	0.070***
	(7.673)	(8.034)	(8.039)	(8.131)
Confirmation bias (γ_S)	0.005	0.003	-0.005	-0.010
	(0.893)	(0.382)	(-0.705)	(-1.322)
Confirmation, Wald $\chi^2(1)$	18.484***	21.524***	25.399***	27.234***
	[0.000]	[0.000]	[0.000]	[0.000]
Herding bias (γ_T)	0.006	0.010	0.015*	0.014*
	(0.854)	(1.293)	(1.931)	(1.660)
Herding bias (γ_S)	-0.014**	-0.019***	-0.023***	-0.019**
	(-2.326)	(-2.973)	(-3.385)	(-2.535)
Herding, Wald $\chi^2(1)$	2.664	4.848**	7.665***	4.824**
	[0.103]	[0.028]	[0.006]	[0.028]

Notes. This table reports the headline specification of Table 3 as the pre-window length L used to construct the behavioural covariates varies across $L \in \{3, 6, 9, 12\}$. The $L = 6$ column reproduces the headline estimates; the other columns use the same identifying-only sample restriction but with covariates averaged over shorter or longer pre-windows. The dependent variable, products, indicators c_T and c_S, definitions of the automation step (γ_T) and the full-delegation step (γ_S), the construction of baseline controls (which are included in every regression but suppressed here), the entity fixed effect on the investor-month, and the two-way clustering on investor and month are as in Table 3. Sample sizes vary with L because the pre-window eligibility filter requires non-missing covariate observations: $L = 3$ uses 226,125 unique investor-months stacked into 678,375 observations; $L = 6$ uses 199,343 stacked into 598,029; $L = 9$ uses 180,508 stacked into 541,524; $L = 12$ uses 151,997 stacked into 455,991. Cluster-robust t-statistics appear in $(\cdot)$ below coefficient estimates; Wald p-values appear in $[\cdot]$ below the corresponding $\chi^2(1)$ statistic. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-tailed).

Table A.9: 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 4 (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 4.

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.10: 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.11: 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.12: 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.13: 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.13 (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.14: 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.15: 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.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.

Table A.16: 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						
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.