Human Financial Advice in the Age of Automation*

Fiona Greig[1] Tarun Ramadorai[2] Alberto G. Rossi[3]

Stephen P. Utkus[3],[4] Ansgar Walther[2]

[1]Vanguard, [2]Imperial College and CEPR, [3]Georgetown University, [4]Wharton, UPenn

July 7, 2025

Abstract

We analyze a unique hybrid robo-advisory setting where portfolio management is automated, but clients are quasi-randomly assigned to human financial advisors who deliver all remaining wealth management services, including behavioral coaching and relationship support. We show that advisors of different types cause variation in client retention, especially during market downturns. To interpret these effects, we estimate a structural model in which advisors influence both investor learning about the automated solution and the perceived utility of remaining enrolled, independent of performance. Effective human advisors generate significant surplus for both clients and the firm through these informational and behavioral channels.

*We thank Brian Martin at Vanguard for useful conversations on institutional details, Paige Weisman for excellent research assistance, and seminar participants at BI Norwegian Business School, Cambridge University Judge Business School, University of Washington Seattle Foster Business School, New York University, Oxford University Said Business School, the University of California at Berkeley, Aalto University, the NBER Behavioral Finance Working group 2023 meeting, the AEA 2023 Structural Behavioral Finance session, and the Swedish House of Finance AI & Machine Learning in Finance conference for comments, and Greg Buchak and Asesh Rambachan for helpful discussions.

Introduction

Technological automation has recently advanced into highly-skilled industries such as financial services, portfolio management, accountancy, and legal services. Recent studies emphasize, for example, the role of machine-based stock market analysis (Cao et al., 2021, 2023), traditionally a task reserved for highly educated human experts.

This paper studies the domain of retail financial advice to analyze whether and how human experts add unique value as financial services become increasingly automated. The possible economic effects in this setting are subtle. On the one hand, the primary task of financial advisors is to recommend financial decisions, such as investment portfolios, to their clients based on their technical skills. For this task, cost-effective, automated "robo-advisors" already offer a popular substitute.[1] On the other hand, human financial advisors also perform auxiliary tasks that require "soft" interpersonal skills such as providing emotional reassurance, which can encourage clients to take on compensated risks (Linnainmaa et al., 2018a), reduce behavioral biases such as loss aversion (Calvet et al., 2023), or increase trust in the financial system (Guiso, Sapienza, and Zingales, 2008; Gennaioli, Shleifer, and Vishny, 2015). While recent research suggests that robo-advice could engender similar behaviors (Linnainmaa et al., 2018b; D'Acunto, Prabhala, and Rossi, 2019), it remains an open question whether there are unique tasks that can be performed only by human advisors in client-facing roles, where interpersonal skills may be particularly important.

Studying the domain of automated financial advice is important because it bears the promise of lowering costs and democratizing access to high-quality financial management (Reher and Sokolinski, 2024). These services have routinely been delivered by humans and have often been accessible only to the wealthy (Gomes, Haliassos, and Ramadorai, 2021; D'Acunto and Rossi, 2022). Concerns have also been raised that human advisors may be prone to biases themselves (Linnainmaa, Melzer, and Previtero, 2021; Andries, Bonelli, and Sraer, 2024) or engage in misconduct (Egan, Matvos, and Seru, 2019).[2] In this sense, automating financial advice could have beneficial implications for both economic efficiency and equity. By contrast, if uniquely human skills are needed to deliver effective financial

[1] Robo-advising has already been widely adopted, which is reflected in its current scale and expected growth—estimated at over $10 trillion under management over the next decade (see, e.g., Deloitte and Touche, "The expansion of robo-advisory in wealth management," 2016).

[2] See Reuter and Schoar (2024) for a useful recent review.

advice, the consequences of automation may be more nuanced, and careful design may be needed to leverage the available technology effectively.

The data on robo-advising that we employ come from a large US robo-hybrid advice service, and span the years 2014-2018. The term "hybrid" refers to the pairing of algorithmic robo-advice with human advisory input.[3] In this service, an algorithm formulated a financial plan and an associated investment portfolio after eliciting investors' characteristics (age, income, investment horizon, and preferences).[4] The sign-up process then required the investor to schedule an appointment with a human financial advisor who explained the financial plan, completed onboarding, and provided ongoing support to investors. The data provide information on trades, positions, demographic characteristics, and investor-advisor interactions for a large set of previously self-directed investors who participated in the robo-advisor.

Critically, the assignment of investors to advisors followed mechanical rules over the period of study. This assignment was driven by workload balancing imperatives rather than any assessment of advisor or client type. Once the current client "load" of a given advisor is accounted for, we confirm that there is no empirical relationship between the historical retention rate of clients assigned to a given advisor and the assignment of new clients to the advisor. This quasi-random assignment allows us to analyze the value of human advice by using advisor fixed effects. We first calculate the rate at which advisors retain clients assigned to them. We then study the causal treatment effects on client behavior of being assigned to a high-client-retention advisor compared to a low-client-retention advisor. We are careful when doing so to avoid any mechanical association by using a "leave-one-out" estimator of advisor type (see, e.g., Collinson et al. (2022)).

Retention rates in the service are generally high, but the data reveal that assignment to advisors of different types clearly predicts retention. Put differently, assignment to low-retention advisors increases the likelihood of exiting robo-advising relative to assignment to high-retention advisors. Moreover, human advice is particularly effective when negative signals of returns hit investors—high-retention advisors' clients are less likely to quit robo-advising during periods of poor performance than

[3]Hybrid robo-advisors allow investors differing levels of interaction with human financial advisors in addition to the algorithmic solution, as opposed to "pure" robo-advisors, which only provide access to the algorithmic solution.

[4]A financial plan principally involves a client-articulated financial goal such as retirement, a cash flow forecast, and a probabilistic assessment of achieving the goal. The portfolio strategy is recommended to implement the plan.

low-retention advisors' clients.

Under the null hypothesis that there is no value to human advice after portfolio choices are automated, assignment to different advisor types would have no predictive power for subsequent client behavior. Our results, therefore, suggest that human advisors do add value. To further interpret these results, recall that the portfolio choice dimension of financial advice is performed by the automated part of the robo-advice service in our data. To corroborate this claim, we verify in the data that the assignment to different types of human advisors has no significant effect on clients' portfolio decisions or realized returns on investment. Therefore, human advisor fixed effects in retention by design do not reflect advisors' technical ability to recommend financial decisions but, rather reflect the auxiliary skills of human advisors for which the automated part of the service does not offer a perfect substitute. An important note is that our detected effects of human soft skills are not necessarily unique to robo-advice, and may also be present in contexts where portfolios are both recommended and implemented by human advisors. This admits an interpretation of our estimates as the *residual* component of human skill, which continues to matter when portfolio choices are delegated to machines.

The remainder of the paper is dedicated to unpacking the economic channels behind the measured treatment effects of human advisor assignment. On the one hand, the observed effects could arise because of the fixed preference effects explored, for example, by Gennaioli, Shleifer, and Vishny (2015): High-retention human advisors may increase clients' trust in the service, their overall willingness to hold risky investments, or the "warm-glow" utility investors perceive from being enrolled and looked after. It would also be reasonable that these effects become more salient for clients' decisions when they have observed disappointing returns, i.e., when they are closer to the margin of quitting the service. On the other hand, the effects could also be driven by information provision. For instance, suppose that clients are initially uncertain about the risk premium that can be earned within robo-advice. If high-retention advisors increase client confidence by providing relevant information, then their clients would become less likely to quit, and differentially less likely after observing poor performance.[5]

Consistent with the idea that clients partly base their decisions on inferred information about risk premia, we show empirically that the effects of human advisor assignments are more pronounced for

[5]Indeed, in standard rational or behavioral models of belief updating, such as the one we outline below, the responsiveness of beliefs to observed *ex-post* signals is smaller when agents have tighter/more confident prior beliefs *ex-ante*.

investors who have had only a short tenure within robo-advice (and have therefore observed only a short time series of returns from which they can make inferences). While it is challenging, in general, to disentangle preference-based and belief-based explanations for agent behavior, we pursue two avenues that help us draw additional inferences about the effects of human advice.

First, we explore the economic channels empirically using client-advisor meeting logs, investor attention behavior, and exit surveys collected from investors who quit robo-advice. Our results suggest that high-retention advisors are effective because they give their clients confidence upfront by meeting with them more often in normal times; consistently, we find that their clients are less likely to quit in times of negative returns. By contrast, low-retention advisors become more active in communicating with their clients only after poor returns have been observed, at which point it is too late to reduce those clients' propensity to quit. High-retention advisors also appear better at communicating the value of the robo-advising services, thus reducing the time investors spend looking for information on the robo-advisor's website. Finally, investors assigned to low-retention advisors more frequently mention wanting to self-manage their portfolio and disagreeing with the robo-advisor's methodology as their main reasons for quitting robo-advice. By contrast, investors assigned to high-retention advisors more often mention the costs of the service as a key driver of their attrition. Investors' actions are consistent with their statements in the exit surveys: those assigned to low-retention advisors are 50% more likely to stay as self-directed investors with the asset manager providing the service even after quitting robo-advice, suggesting that their dissatisfaction is with the service rather than the asset manager.

Second, to more deeply understand the underlying sources of advisor effectiveness, and to enable quantification and counterfactual analysis, we build a structural model of human financial advice in an environment of automated portfolio solutions. We analyze the dynamic choice problem of an investor who is initially enrolled in hybrid robo-advice, draws Bayesian inferences about the returns provided by the service, and derives a stream of ongoing utility or disutility from enrolment that is independent of wealth accumulation. The investor decides in each period whether to remain enrolled or to quit the service, aligning with our empirical work that studies investor attrition rates. The investor in the model is matched with a human advisor of either low or high effectiveness, to capture the variation in

retention across advisors observed in our empirical analysis. The advisor can affect both the rate of learning about the returns from the service (through both the prior mean and prior precision of beliefs about these returns), as well as the ongoing benefits from enrolment. The effects of advisors on the ongoing benefit stream flexibly capture the use of interpersonal skills to alleviate algorithm aversion, distrust of automated solutions, greater trust, or anticipation of higher risk-adjusted returns. Our specification of advisor effects on prior beliefs also admits both rational and behavioral interpretations. A rational interpretation of the positive effects of high-effectiveness advisors on prior mean returns and prior precision of returns is that such advisors communicate information about the service in a more precise/less noisy manner than others. An alternative, behavioral explanation is that human advisors can affect client beliefs through coaching, handholding, or persuasion.

We structurally estimate the parameters of the model to match non-parametrically estimated hazards of investor attrition from the service. We do so for investors matched to both high- and low-retention advisors, and those who have experienced low and high average returns during their tenure. The estimates reveal two dimensions of the structural differences between high- and low-retention human advisors. First, we find that the precision induced by a high-retention advisor is equivalent to observing a time series of roughly 20 years of performance data, at the start of enrolment, whereas the equivalent for a low-retention advisor is equivalent to roughly 12 years of performance data. This result on the belief channel is consistent with high-retention human advisors adding value by increasing their clients' confidence in product quality. Second, we find that high-retention advisors deliver greater flow utility to investors, translating to a lifetime benefit that is equivalent to a portfolio return of 0.3% return (30bps) per annum.[6]

We use our parameter values to quantify the surplus obtained by investors, as well as by the advisory firm, in different advice scenarios. In money-metric terms, we find that a lump-sum transfer equal to 9.6% of her initial wealth would make the investor indifferent between being matched with a low-retention and a high-retention advisor. From another perspective, we find that the contribution to the firm's surplus from the lower quit rates induced by high-retention advisors is around 0.26% (26bps) per dollar of investor assets, a significant effect given the scale of major financial advisory

[6]Coincidentally, this enhanced benefit is roughly equivalent to the fees for the service, admitting the interpretation that the service "feels free" to investors matched with high-retention advisors.

firms. These gains mainly come from the increase in ongoing flow utility benefit of high-retention advisors; while the gain in confidence appears economically significant, it is less important in welfare terms. One interpretation of this finding is that adroit advisors help mainly through alleviating distrust of automation. Another is that investors derive substantial additional "warm glow" utility à la Gennaioli, Shleifer, and Vishny (2015) from interactions with their advisors.

Our work contributes to several strands of the literature. First, we contribute to the growing literature on robo-advising (in addition to previously referenced papers, see, for example, Reher and Sun (2019); Rossi and Utkus (2019, 2024)). This literature is now beginning to explore the role of trust in robo-advice, linked to the broader literature on the role of trust in financial institutions and stock market participation (see, e.g., Guiso, Sapienza, and Zingales (2008)). Second, we contribute to the household finance literature, which seeks ways to both incentivize households to participate in financial markets and do so in an efficient, well-diversified manner (see, e.g., Campbell (2006); Guiso and Sodini (2013); Badarinza, Campbell, and Ramadorai (2016)). Third, our work is related to the more general literature on the quality of financial advice (see, e.g., Linnainmaa et al. (2018a); Egan, Matvos, and Seru (2019). In this setting, our focus is primarily on the ability of advisers to assuage investors' concerns about the algorithmic solution.

The remainder of this paper is organized as follows: Section 1 describes the institutional details of our empirical setting and the details of our identification strategy. Section 2 contains our main empirical results, Section 3 contains our structural modeling exercise, and Section 4 concludes.

1 Institutional Setting and Data

We use data from a large US-based hybrid robo-advisor. The service is "hybrid" as it complements automated portfolio allocations with human advisory input. More broadly, robo-advisors are commonly classified as either pure or hybrid robo-advisors. Pure robo-advisors do not feature any substantive interactions between investors and human financial advisors, whereas hybrid robo-advisors allow investors differing levels of interaction with human financial advisors. In our setting, the investment portfolio management was automated, while the human adviser interacted with investors to help them understand these automated investment decisions and how they facilitate clients' financial goals, as

well as advising on related services such as tax strategies and debt management.

1.1 Characteristics of Advised Investors

Our data run from 2014 to 2018, and contain information on trades, positions, demographic characteristics, and investor-human advisor mappings for previously self-directed investors who signed up for (or considered signing up for) the robo-advisor service. The trade data includes all trades placed by the investors when self-directed and by the robo-advisor when the investor is in the service. The position data consists of associated monthly holdings observations, tracking both investments and trades. Additionally, the data record the dates when investors initiated, enrolled, implemented, and quit the advice service. Demographic characteristics include investor age, gender, and tenure with the asset manager. The investor-advisor mapping data tracks the dates, times, lengths, and initiators (advisor or investor) of all advisor-investor interactions, including meetings and phone or video calls.

Table OA.I reports summary statistics (mean, standard deviation, and percentiles of the distribution) of different variables characterizing advised investors in the sample, and recorded 12 months after investors signed up for the service. Panel A focuses on investors' demographic characteristics. The average investor is 64 years old, 60% of the users are male, and the average investor has had an account with the asset manager for almost 14 years. The advised population comprises older, wealthier investors who are more gender-balanced than datasets commonly used to study trading behavior.[7] Table OA.I Panel B shows details of investors' portfolio allocations. The average investor has $758,378 invested with the robo-advisor, which is more than 50% larger than the median ($478,929), reflecting significant right-skewness. There are 8 distinct assets on average in each account, comprising mutual funds, stocks, bonds, and ETFs, with the bulk of the investors being invested in only 6 assets, and only 25% of the investors having more than 9 assets in their portfolio. This reflects the different "glide paths" to which investors are assigned by the automated service (i.e., their target asset allocations differ based on age, wealth, and other characteristics); we discuss this in greater detail below. Almost all investor wealth (97.4%) is invested in products provided by the robo-advisor's asset manager (mostly indexed mutual funds). Finally, Table OA.I Panel C shows that 95% of investors' wealth is in mutual

[7] In contrast with the demographics in these data, the average investor age is 51 in the brokerage trading data employed by both Gargano and Rossi (2017) and Barber and Odean (2001), and women constitute lower fractions in those datasets, with 27% in Gargano and Rossi (2017) and 21% in Barber and Odean (2001).

funds, followed by cash (2%), individual stocks (1.4%), and ETFs (0.8%). Only a negligible number of investors have direct exposure to corporate bonds and options (not reported).

Given that the majority of investors' wealth is in mutual funds, in Panel D, we analyze the characteristics of the mutual funds held. The first row shows that 83% of mutual fund holdings are in indexed mutual funds. The average management fee across all mutual fund holdings is 7 bp. The expense ratios are also low, on average 9 bp with a median of 8 bp. Finally, the average turnover ratio of the mutual funds in the portfolio is 27%.

1.2 Robo-advisor Characteristics and Returns of Advised Clients

During the sample period, the robo-advisor classified investors into five risk glide paths based on their financial objectives, investment horizons, and demographic characteristics. While confidentiality prevents us from disclosing the specifics of the algorithm used to generate the portfolio allocations, we present key statistics on the general characteristics of these allocations and the performance of the advised portfolios.

We compute monthly portfolio returns for all advised investors; these investors are categorized into two tiers based on their investment commitment. Higher-tier investors have a dedicated investment advisor, while lower-tier investors receive support from a pool of human advisors.[8] For the purposes of these summary statistics, we co-mingle the two groups because their investment performance is virtually identical. For comparison, we also compute the average returns for self-directed clients who interacted with the robo-advisor during the sample. This includes the returns of clients who later signed up for robo-advice, as well as those who did not subsequently sign up after considering the service. Figure 1(a) shows the distributions of (net-of-advisory fee) monthly returns for both advised investors (blue) and self-directed investors (red), computed across all periods and investors. The plot reveals that advised investors achieve higher mean returns than their self-directed counterparts, as indicated by the rightward shift of the blue distribution relative to the red distribution. The plot also reveals lower cross-sectional variation in returns for advised investors, likely attributable to the better diversification of their portfolios.[9] We provide additional details on the returns distribution of

[8]Section 1.4 describes the (quasi-random) process of assignment of dedicated human advisors to higher-tier clients; lower-tier investors receive support based on scheduling/availability from the pool.

[9]The red distribution shows notable bunching at zero due to some self-directed investors holding their entire portfolios

self-directed and advised investors when we calibrate and estimate our structural model in Section 3.

To better understand the sources of variation in the returns of advised investors, we conduct the following three-step procedure. In the first step, following Balasubramaniam et al. (2022), we compute principal components (PCs) of the equity share of 1,700 investors for which we have 36 continuous months of advised returns. The first PC explains 46% of the variation in equity shares, and the first five PCs explain 81% of this variation. These results are the first indication of a pronounced factor structure in the robo-advisor portfolio allocations, arising from the small number of glide paths to which investors are assigned. In the second step, we regress the time series of the equity share for each of the 55,000 investors in our data on the first five PCs estimated in the first step. Finally, in the third step, we cluster investors' loadings in five groups using a k-nearest neighbors estimation procedure. We find that 82.4% of the investors are classified in the first group, followed by the second and third groups with 13.5% and 3.4% of the investors, respectively. Finally, the last two groups contain less than 1% of the advised investors. All in all, these results suggest that, while five risk glide paths exist, the majority of investors end up on two glide paths over our sample period.

We complement this simple analysis of equity shares by regressing monthly investor total portfolio returns on the market portfolio, investors' equity share, and the interaction of these two variables. This regression has an R^2 of over 75%, highlighting that the equity-bond allocation decision and the variation in aggregate equity returns together explain the major share of variation in returns seen in the data. We report realized monthly returns against the predicted returns from this regression in a binned scatterplot, Figure 1(b), which shows that the two quantities line up very closely along the 45° line.

Given the robo-advising glide paths, the cross-sectional variation seen in the equity share and returns is likely an outcome of variation in investors' preferences and demographic characteristics, which the automated service translates into different portfolio allocations. While we do not have information regarding investors' risk preferences, we do have investors' age. We therefore regress monthly client returns on the market portfolio, investors' age, and their interaction. This regression has an R^2 of over 71%, which is very close to the R^2 obtained using equity shares. The corresponding binned scatterplot of realized vs. fitted returns reported in Figure 1(c) suggests that any risk preference

in cash during certain months.

variation over and above that which is correlated with age likely plays a minor role in determining the portfolio allocation implemented by the algorithm.

Figure 1(d) provides a different perspective. The black line shows the distribution of cross-sectionally demeaned average annualized returns on all advised portfolios. The red and blue solid lines show the residuals from regressions on the market and equity share, and the market and age, respectively. The substantial reduction in the variance of the distribution of these residuals shows that equity share and age explain a majority of the variation in average portfolio returns, which means that performance is, in large part, homogenous for investors with similar demographic characteristics.

Figure 1(d) also shows dashed red and dashed blue distributions, which include advisor fixed effects into the regressions of performance on equity share and age, respectively. If advisors were instrumental in affecting the portfolio allocations recommended by the algorithm, we would expect these two distributions to have a smaller variance than the corresponding solid distributions of residuals from models that do not include advisor fixed effects. Instead, we find that the dashed blue and red distributions are virtually identical to the respective solid distributions, suggesting that individual advisors play virtually no role in affecting investors' portfolio allocation. In the remainder of the paper, we use human advisor-client interactions to assess the degree to which human advisors are complementary to the automated portfolio strategy. Along with the random assignment of human advisors to clients, Figure 1(d) provides evidence that this complementarity in the data arises from sources other than human advisors directly affecting portfolio strategy, aiding parameter estimation from the structural model.

1.3 Measuring Advisor Type

We seek to understand the effects of human financial advice in this setting. In the absence of an experiment that assigns some clients randomly to human advisors while leaving others without human input, we concentrate on measuring the types of human financial advisors. Measuring advisor type is challenging for several reasons. First, success in this setting is likely the combination of many different but complementary traits, including client-specific assessments of advisor communication and relationship-management skills. Second, the skills needed to be a successful "plain vanilla" financial

advisor may be different from those needed as a hybrid robo-financial advisor, so it is difficult to rely on "standard" personality assessments or financial competence scores in this novel setting. Third, an important determinant of customer satisfaction in the financial domain is portfolio performance, which is determined by the automated algorithm, not the advisor. As a result, performance-based metrics—a common approach in financial economics—are not appropriate for evaluating advisor type in this context.

We therefore rely on revealed preference and measure advisor type using advisor-client-specific average retention rates. More specifically, we construct a measure of advisor retention for each investor using a leave-out estimator. This approach follows the logic of Collinson et al. (2022), who construct judge-specific leniency measures by excluding the focal observation to avoid endogenous correlation between the instrument and the outcome. It eliminates the mechanical correlation between the treatment variable (advisor type) and the outcome (retention), which would otherwise arise if an investor's own outcome were used to compute the type of advisor to whom they are assigned.

For each client, we compute the average client retention of the human advisor they are assigned to, using all clients assigned to the advisor except for the client in question. Specifically, the retention rate of advisor j, applied to investor i, is the ratio (estimated over the full sample) of all clients of advisor j (excluding investor i) who remain in the robo-advice service, divided by the total number of clients assigned to advisor j. In our main specifications, if an investor i quits, we exclude the period in which they quit to eliminate any potential bias arising from contemporaneous correlation in quits across investors assigned to the same advisor. We then discretize the retention measure, categorizing advisors as high retention or low retention using the median retention rate as the breakpoint.[10] As a first step, we use advisor-specific client retention to measure advisors' type, without taking a stance on specific skills that increase retention. We then analyze the sources of these skills using further empirical results on advisor-client communications, as well as the structural modeling exercise in Section 3.

[10]We note that the overall retention rate in robo-advising is high—above 95%—across all advisors, we detail this further in Section 2.

1.4 Assignment of New Investors to Advisors

During the sample period, the robo-advisor assignment and onboarding process began by eliciting investors' characteristics (age, income, wealth, investment horizon, and preferences). Using this information, the robo-component of the robo-advisor (i.e., the algorithm) automatically formulated a financial plan and an associated investment portfolio. The sign-up process then required the investor to schedule a time to meet with a human financial advisor who explained the financial plan and completed onboarding.

The assignment of clients to advisors was quasi-random. Advisors were assigned based on the match between a prospective client's availability for an initial onboarding call and the availability of advisors in the system. Advisor availability was driven by the need to balance workload across advisors (i.e., driven by advisor capacity) rather than any assessment of advisor type (i.e., advisor skills, characteristics, or ability to retain clients). Put differently, the assignment of new clients to advisors during the sample period was random with respect to advisor type, depending on advisors' capacity and availability but not their skills or characteristics.

To verify random assignment, Table 1 shows the main attributes of investors assigned to high- and low-attrition advisors the month before they sign up for robo-advice. The rows are separated into different sets of covariates: 1) demographic characteristics of the investors, such as age, gender, and the tenure of investors with the asset manager before signing up for the robo-advisor service; 2) portfolio-related characteristics, such as the total assets held by the investor at the asset manager, the number of assets, and the percentage of the asset manager's products (mutual funds and ETFs) in investors' portfolios; 3) investors' asset allocation characteristics, such as the percentage of assets in mutual funds, cash, ETFs, individual stocks and bonds at the time of sign up; 4) characteristics of the mutual funds held by the investors, such as degree of indexation, fees charged, expense ratios and turnover ratios; and 5) investment performance measures, such as the average annualized monthly returns realized by the investors before signing up for robo-advice as well as the month-by-month cross-sectionally demeaned version that controls for time-variation in returns.

Across all characteristics, the covariates are balanced in that investors assigned to high- and low-attrition advisors are similar in their portfolio sizes, portfolio allocation across asset classes, the types

of mutual funds they invest in, and their returns before signing up for the service. The only covariates where we detect statistical differences are age, gender, and tenure at the asset manager before signing up. In all cases, however, even when the differences are statistically significant, they are economically small. For example, the average age of those assigned to high-retention advisors is 64.5, while it equals 65.6 for those assigned to low-retention advisors.[11]

Figure OA.V shows the relationship between advisor load, advisor leave-out retention, and the assignment of clients to advisors. For each advisor, we compute the number of clients they advise at the beginning of each month and sort advisors into quintiles based on their current workload. Figure OA.V (a) shows the net increase in the number of clients allocated to advisors in each group on average every month, computed as the number of investors allocated to each advisor minus the number of investors lost by each advisor every month due to attrition, together with 95% confidence intervals. The plot shows that there is a ramp-up in the rate of net additions from the first up to the third quintile of current capacity and a decline from then to the fifth quintile as advisors reach full capacity. The top group of advisors that are close to full capacity has a monthly net average growth that is not statistically different from zero. Figure OA.V (b) repeats the analysis sorted using deciles of current capacity instead of quintiles, with qualitatively similar results.

For valid causal inference about the treatment effects of advisor assignment, we need the assignment of clients to advisors to be independent of performance and preference shocks. Table 1 demonstrates that client characteristics are balanced across advisor types. Figures OA.V (c) and (d) plot client assignment rates to advisors of different types. These figures show gross additions to advisors (i.e., ignoring attrition) and split advisors into high- and low-retention, conditional on current capacity ((c) splits into quintiles of current capacity and (d) deciles). Except for the very first group in both plots, where we find an economically small difference in the assignment rates to high-retention advisors (in blue) and low-retention advisors (in red), all other capacity groups show no differential rates of assignment to advisors with different retention rates.[12]

[11] Section 2.3 demonstrates that our results are robust to controlling for the full set of observable investor characteristics.

[12] We later eliminate the lowest current capacity group in our results and verify robustness. We note that with ten deciles, finding one statistically significant difference is consistent with a 10% rate of significance.

2 Empirical Estimates

The random assignment of advisors to clients allows us to estimate the causal effect of human advisors of different types on retaining investors in the robo-advice service.

2.1 Cross-sectional Variation in Advisor Retention

Figure 2 (a) shows a histogram of the (leave-one-out estimated) advisor retention measures, scaled in such a way that the highest leave-out retention estimates in the data are set to 100. The figure shows substantial cross-advisor dispersion in the retention measure. Superimposed on the histogram in red is a non-parametric (lowess) estimate of the retention rate of clients who are randomly assigned to advisors of different quality, measured by their scaled leave-one-out retention rates. The plot shows a clear near-linear relationship between the retention rate of clients that are randomly assigned to advisors and the quality of these advisors.[13] One possible concern with Figure 2 (a) is that quality may be mismeasured because of outliers caused by small base effects—i.e., advisors who never got to manage a large set of clients. Figure 2 (b) therefore removes advisors in the bottom decile of clients under management throughout the sample period; and shows that the picture is virtually unchanged.

2.2 Baseline Non-Parametric Survival Estimates

We next estimate separately Kaplan-Meier survival functions for clients assigned to low- and high-retention advisors:

$$\widehat{S}(t) = \Pi_{k:t_k \leq t} \left(1 - \frac{d_k}{n_k}\right), \tag{1}$$

where t_k is a time when at least one investor quits robo-advice, d_k is the number of clients quitting robo-advice at time t, and n_k is the number of individuals who remained enrolled (i.e., neither quit nor were censored) up to time t_k.

Figure 3 (a) plots Kaplan-Meier survival functions by advisor type. The blue line (and associated 95% confidence intervals) shows survival rates for clients who are assigned to high-retention advisors,

[13] This figure is similar to Figure 5 in Collinson et al. (2022) who use a judge stringency instrument.

and the red line shows survival rates for investors assigned to low-retention advisors. The figure shows that 96.2% of investors in the high-retention advisor group and 95.5% of investors in the low-retention advisor group stay with robo-advising at the one-year mark. Over time, the survival gap widens: at the three-year mark, 90.6% of clients assigned to high-retention advisors are still enrolled, compared to 86.8% for those assigned to low-retention advisors.

Figure OA.V shows that high- and low-retention advisors are assigned very similar numbers of clients, conditional on their current load. The only exception is the bottom decile—the advisors with the lightest load—where high-retention advisors are assigned more investors: 1.5 on average per month versus 0.8. To ensure that these advisors do not drive the results, Figure 3 (b) reports survival estimates that exclude this bottom decile; the results are virtually identical.

As a complementary approach, we estimate a Cox proportional hazard model for attrition from the robo-advising service. Figure 3 (c) plots smoothed hazard functions to assess whether the proportional hazards assumption holds. The hazard rates are approximately parallel, with greater attrition for clients assigned to low-retention advisors. For these clients, the hazard estimates (red line and confidence interval) increase and peak at 6% at the one-year mark, only to decline and reach their bottom of 2.8% at the three-year mark. The blue line (and confidence interval) shows hazard estimates for those assigned to high-retention advisors, which follow similar dynamics at lower levels: the hazard estimates peak at 4.2% at the one-year mark, declining and bottoming out at 1.9% in year three. Figure 3 (d) replicates the analysis excluding the bottom decile of advisors, again yielding virtually identical results.

2.3 Robustness and Validation

To strengthen the interpretation of our baseline findings, we address several potential concerns. Specifically, we examine whether observable client characteristics, time variation, model specification choices, the reassignment of clients to different advisors, or unobserved onboarding dynamics could account for the differential attrition rates across advisor types.

First, while Kaplan-Meier survival estimates indicate that advisor type predicts client retention, they do not condition on investor characteristics. Although Table 1 shows that covariates are generally

balanced across advisor types, small imbalances could still contribute to the observed differences in attrition.

To address this concern, we estimate a Cox proportional-hazard model of the following form on all advisors:

$$h(t|x_j) = h_0(t) \cdot \exp(x_j\beta).$$

where x_j denotes advisor type and investor covariates. We first include only the "high-retention" advisor indicator as a regressor. We then sequentially add investor covariates from Table 1, ordered by the absolute value of their t-statistics—from largest to smallest.

Figure 4 (a) plots the estimated coefficient on the high-retention indicator and its 95% confidence interval. Without controls, the coefficient β equals -0.293, implying that clients assigned to high-retention advisors are $\exp(-0.293) = 0.746$ times as likely to quit—a 25% reduction in hazard. Adding covariates has a negligible impact on the coefficient magnitude or significance, indicating that observable client characteristics do not explain the differential attrition.

Second, our advisor classification does not explicitly control for time variation. If some advisors were disproportionately active during periods with higher baseline quit rates, time effects rather than advisor quality could drive the results. To test this, Figure 4 (b) replicates the analysis while adding increasingly granular time fixed effects (from annual to weekly). The coefficient estimates and their confidence intervals remain stable, suggesting time effects do not explain the findings.

Third, discretizing advisors into only two groups may obscure important variation across the distribution of advisor quality. To address this, Figures 4 (c) and (d) replace the high-retention dummy with a continuous measure of advisor retention. Without controls, the coefficient $\beta = -0.045$ implies that a 1% increase in advisor retention is associated with a 4.4% reduction in the hazard of quitting ($\exp(-0.045) = 0.956$). Controlling for client characteristics and adding time fixed effects does not materially change the estimates.

Fourth, our baseline analysis assigns each client to their initial advisor when computing results. However, due to advisor retirements or moves to other firms, some clients may be re-assigned and interact with multiple advisors over time, potentially confounding the relationship between advisor

type and client attrition. To test for this, we re-estimate our main results restricting the sample to clients who interacted with only one advisor during their tenure with the robo-advisor. The results, reported in Figure OA.VI, show that the estimates are virtually unchanged.

Finally, one potential concern is that differences in client retention could be driven by the onboarding process itself. In particular, if low-retention advisors were more effective at attracting new clients during the initial meeting, they might enroll clients who are less committed to the service and thus more likely to quit later. However, if this mechanism were operative, we would expect low-retention advisors to systematically add more new clients over time compared to high-retention advisors, conditional on their current workload. In contrast, Figures OA.V (c) and (d) show that client assignment and net client growth rates are similar across advisor types. This evidence suggests that onboarding-based selection is unlikely to drive the differential attrition we document.

Taken together, these robustness exercises strengthen the interpretation that assignment to higher-retention advisors is associated with greater client retention, consistent with a causal role of human advisor type in maintaining engagement with robo-advising.

2.4 Estimating the Effect of Human Advice Across Investment Performance

We next evaluate how client quit rates vary with investment performance and advisor type. To this end, we estimate the following regression at the monthly frequency:

$$\begin{aligned} Dummy_quit_{i,t} = \alpha + \beta\ I_{\{Negative\ \ Return_{t-1}=1\}} + \gamma\ I_{\{HighRetn_Advisor_i=1\}} \\ + \delta\ I_{\{Negative\ \ Return_{t-1}=1\}} \times I_{\{HighRetn_Advisor_i=1\}} + \epsilon_{i,t}, \end{aligned} \tag{2}$$

where $Dummy_quit_{i,t}$ is equal to 1 if investor i quits robo-advising in month t and 0 otherwise, $I_{\{Negative\ \ Return_{t-1}=1\}}$ is an indicator variable equal to 1 if the investors' lagged returns are negative in month $t-1$ and 0 otherwise, and $I_{\{HighRetn_Advisor_i=1\}}$ is equal to 1 if investor i is assigned to a high-retention advisor and 0 otherwise. We also estimate a version of equation (2) that includes only the $HighRetn_Advisor$ dummy.

In equation (2), α measures the monthly quit probability for investors assigned to low-retention advisors in periods in which their portfolio has performed well, β estimates the conditional increase

in attrition from robo-advice in periods in which their portfolio has performed poorly, γ measures the differential quit rate associated with clients assigned to high-retention advisors, and δ measures the extent to which the differential quit rate across clients assigned to high- and low-retention advisors varies with market conditions. In the version that includes only the $HighRetn_Advisor$ dummy, α measures the monthly quit probability for investors assigned to low-retention advisors, and γ measures the differential quit rate associated with clients assigned to high-retention advisors.

Table 2 presents the estimates of this equation, with all coefficient estimates scaled by a factor of 100 to express magnitudes in percentage points. The first column indicates that investors assigned to low-retention advisors have a monthly quit rate of 0.473% (equivalent to 5.68% annually), in line with Figure 3. In addition, the γ coefficient shows that being assigned to a high-retention advisor results in a monthly quit rate that is 0.127 percentage points lower (or 1.52% less per year), confirming the results from our survival analysis. In the second column, we use investors' lagged average returns since enrollment into robo-advice to capture variation in portfolio performance. The α coefficient in this column reveals an unconditional 0.442% quit probability per month, roughly translating to an annual attrition rate of $0.442\% \times 12 = 5.304\%$. The β coefficient shows that attrition from robo-advising increases by 0.178 percentage points when investors face poor performance, i.e., when their average returns are lower than zero. This increase is large, translating to a $0.178/0.442 = 40\%$ increase in attrition in such times. The γ coefficient shows that being assigned to high-retention advisors reduces attrition by 0.107% per month. Finally, δ is negative and significant, showing that clients randomly assigned to high-retention advisors have lower attrition than those assigned to low-retention advisors when portfolio performance is low.

In the third column of Table 2, we repeat the analysis using investors' lagged one-period returns rather than their lagged average returns since enrollment. Overall, the results are very similar when we work with this alternative specification: worse portfolio returns are associated with higher attrition rates, high-retention advisors are associated with lower attrition rates, and the effect of an advisor is particularly pronounced when portfolio returns are negative. In Figure 5, we use the regression estimates from Table 2 to plot attrition across different portfolio returns and advisor types. The first set of results is associated with the estimates in the first column of Table 2 and do not condition on

portfolio performance. They show that investors assigned to high-retention advisors have a monthly quit rate of 0.346% (or 4.15% per year), while investors assigned to low-retention advisors have a monthly quit rate of 0.473% (or 5.68% per year). The second and third sets of results use lagged average returns since enrollment to measure portfolio performance. Investors assigned to high-retention advisors have an attrition rate equal to 0.33% per month (3.96% when annualized), which increases to 0.43% (5.16% when annualized) for negative average returns, showing that poor portfolio performance increases attrition by 30%. Investors assigned to low-retention advisors, on the other hand, have an attrition rate of 0.44% per month when their portfolio return is on average positive and a probability of attrition of 0.62% when it is negative, for a percentage increase in attrition equal to 41%. The fourth and fifth sets of results are associated with the estimates in the third column of Table 2 and use lagged one-period returns as a measure of performance. Investors assigned to high-retention advisors have an attrition rate of slightly below 0.35% per month, and this estimate does not vary with lagged one-period returns. Investors assigned to low-retention advisors, on the other hand, have an attrition rate of 0.45% per month when their prior return is positive and a probability of attrition of 0.58% when it is negative. These results are consistent with high-retention advisors reducing attrition both conditionally and unconditionally, i.e., high-retention advisors are particularly important when portfolio performance is poor.

Table 2 columns 4 and 5 further condition the results on investors' attrition-rate sensitivity to performance by looking at this separately for investors that have been with the robo-advice service for shorter (column 4) and longer (column 5) tenures (we operationalize this by splitting above and below the median tenure in the service). The table focuses on recent signals, i.e., one-period lagged returns, and shows that while the unconditional effects of a high-retention advisor are roughly similar for both groups, the interaction effect of having a high-retention advisor when returns disappoint is stronger for long-tenure investors. Put differently, long-tenure investors with high-retention advisors quit far less in response to recent negative returns than shorter-tenure investors. This suggests the presence of learning effects—longer-tenure investors matched to high-retention advisors learn little from recent negative-return realizations as these minor blips are now more familiar to them.

2.5 Explaining Heterogeneity in Advisor Retention Rates

The results thus far demonstrate that advisor type causally impacts investors' retention in robo-advice. We now empirically explore the economic channels of these results using client-advisor meeting logs, investor attention behavior, and exit surveys collected from investors who quit robo-advice.

2.5.1 Client-Advisor Meeting Logs

Our data contains a log of all the interactions between investors and their assigned advisors, including the purpose of each meeting, its duration, and whether the advisor or the client scheduled it.

We first compute, for every investor i, the average number of minutes per year they spend with their advisor over the period they are advised. In the second step, we average this quantity across those investors randomly assigned to high-retention advisors and those assigned to low-retention advisors. The results, reported in the first set of bars (denoted by "Total") in Figure 6, show significant differences in the time investors spend with their advisors. High-retention advisors (blue bar) meet with their clients 133 minutes per year, 15.7% more than those investors assigned to low-retention advisors (red bar), who instead meet 115 minutes per year with their clients. The 95% confidence bands confirm that the two quantities are statistically different from each other. One possible driver of this finding is that customers assigned to high-retention advisors opt to engage more frequently with their advisors due to their satisfaction with the provided services. Alternatively, it could be the high-retention advisors who schedule more meetings with their clients to make sure they are happy with the services they are receiving and understand how the service operates.

The second and third sets of bars in Figure 6 distinguish between advisor-scheduled and client-scheduled meetings. The vast majority of the meetings are advisor-scheduled. Moreover, while we observe an economically significant difference between the time high-retention advisors spend with their clients in the context of advisor-scheduled meetings, we do not find differences for client-scheduled meetings. These results suggest that the channel through which high-retention advisors retain customers is by scheduling meetings and cultivating stronger relationships. The results thus far focus on all client meetings. In the fourth and fifth sets of bars in Figure 6, we further decompose meetings into non-recurring and recurring ones—where the latter include only quarterly, semi-annual, and annual

check-in meetings, which the robo-advisor company mandates. The vast majority of the time advisors spend with clients is due to non-recurring meetings, suggesting that low-retention advisors satisfy their basic client contact requirements. In contrast, high-retention advisors exert additional effort, scheduling more appointments with their clients to further their relationships.

2.5.2 Conditional Results

The client-advisor meeting results reported so far do not distinguish between up and down markets. However, many of our results show that investors assigned to low-retention advisors quit disproportionately more often during periods of market distress. In the last two sets of bars in Figure 6, we test whether the assignment to low- and high-retention advisors affects client-advisor meetings depending on market conditions.

The results for good market conditions align with the unconditional results, and show that investors assigned to low-retention advisors interact *less* with their advisor compared to those assigned to high-retention advisors. The results for bad market conditions show instead that clients assigned to low-retention advisors interact *more* with their advisors in periods of market distress. These findings are consistent with high-retention advisors being effective because they give their clients confidence upfront, and their clients do not panic in down markets. On the other hand, low-retention advisors nurture the relationship with their clients during good market conditions to a lesser extent than high-retention advisors. In turn, during market downturns, they face more jitters from (and have to provide more assurance to) their clients. This extra effort, however, isn't enough to convince all of their clients to stay advised, helping to explain why we observe differential attrition across high- and low-retention advisors in up and down markets.

2.5.3 Login and Attention Information

Next, we use investors' attention and login information to evaluate the channels driving our main findings. In the first two sets of bars in Figure 7, we display the number of minutes individuals spend every year acquiring and processing information from the robo-advisor. We split the overall attention into the minutes spent on the website and the minutes spent on the app.

Unlike the results for client-advisor meetings, investors assigned to low-retention advisors (red bar) spend more time on the robo-advisor's website compared to those assigned to high-retention advisors (blue bar). The differences in attention on the robo-advisor's website are economically large, in that clients assigned to low-retention advisors spend 467/401=16.5% more time on the platform compared to those assigned to high-retention advisors. However, we do not detect differential attention between those investors assigned to low- and high-retention advisors when it comes to the time they spend on the robo-advisor's app. It appears that investors use the app periodically to track the value of their robo portfolio, but not to acquire information about the robo-advisor's operations.

The last two sets of bars in Figure 7 show that clients assigned to low-retention advisors spend more time on the robo-advisor's website in both good and poor market conditions. Moreover, investors assigned to low-retention advisors log in disproportionately more often in down markets, consistent with investors logging in during periods of market distress for information regarding the actions of the robo-advisor.[14]

2.5.4 Exit Surveys

For more than 95% of the investors who discontinued robo-advice, we have access to their exit interview responses. After excluding natural deaths, we count 43 distinct reasons for leaving robo-advice. The three most frequent reasons, which comprise close to 45% of the responses, are 1) the client wants to self-manage (29%); 2) the client has transferred assets out of the asset manager (8%); and 3) the client disagrees with the investment philosophy (6%).

Because the second reason for quitting is not particularly informative, we focus on the first and third reasons and condition the results on advisor type. The left subfigure of Figure OA.VII shows that investors assigned to low-retention advisors disproportionately mention a) disagreeing with the robo-advisor's investment philosophy, and b) wanting to self-manage as the main reasons for quitting robo-advice, compared to investors assigned to high-retention advisors.

In a second exercise, we take all 43 reasons for quitting robo-advice, and classify them into reasons related to investors wanting to self-manage their portfolio and reasons related to robo-advisor's costs

[14]There is also a meaningful "ostrich effect" in logins among all investors, regardless of adviser type—logins in up markets are well above logins in down markets (Sicherman et al., 2016)

and performance. In the first category, we include—among others— reasons like "Wants to self manage" and "Client disagrees with investment philosophy." In the second, we include—among others—reasons like "Too expensive" and "Client is unhappy with portfolio returns." We exclude all other reasons from the analysis. We report the results in the middle subfigure of Figure OA.VII, where we show that investors assigned to low-retention advisors disproportionately mention being unconvinced by the robo-advisor's methods as the main reason for quitting robo-advice. The investors assigned to high-retention advisors, on the other hand, disproportionately mention the costs and performance of the robo-advisor as reasons for quitting robo-advice.

Finally, in the right subfigure of Figure OA.VII, we verify that investors' actions are consistent with their statements in the exit surveys. In particular, we plot the percentage of investors who after quitting robo-advice, maintain their account with the asset manager as self-directed investors. Whether we focus on all investors or only on those who stay self-directed for at least 6 months after quitting robo-advice, our results show that investors assigned to low-retention advisors are more likely to stay as self-directed investors after quitting robo-advice, suggesting that their dissatisfaction was directed towards the robo-advising service and not the asset manager as a whole and that their advisor did not manage to successfully communicate to them the value of robo-advice. These results reinforce the view that by not interacting enough with their clients, low-retention advisors may not have been as successful in explaining to investors how the robo-advisor operates, and this may be the primary channel driving the differences in attrition documented earlier.

For a deeper understanding of the underlying economics of advisor effectiveness, and to enable quantification and counterfactual analysis, we build a structural model of human financial advice in an environment of automated portfolio solutions. We now turn to describing this model.

3 Structural Model

3.1 Model Environment

We model the dynamic choice problem of an investor who is initially enrolled in hybrid robo-advice. Over time, the investor decides whether to remain enrolled or, alternatively, to quit the service.

Time is discrete and indexed by $t \in \{0, 1, \dots\}$. At an initial date $t = 0$, the investor is born with

exogenous wealth W_0 and is automatically enrolled in a hybrid robo-advice service. She is also matched with a human advisor of type $j \in \{0, 1\}$, where j denotes the equivalent of low- and high-retention advisor types in our empirical analysis.

At every subsequent date $t = 1, 2, \ldots$, the investor decides whether to remain enrolled, denoted $z_t = 1$, or to quit, denoted $z_t = 0$. It is not possible to re-enrol after quitting, so that $z_t = 0$ implies $z_s = 0$ for all $s \geq t$.[15] Moreover, at each date, the investor's investment horizon ends with an exogenous and constant probability ρ, at which point she consumes her terminal wealth. We denote the stochastic date of consumption by $t = \tau$.

The investor's preferences are represented by the following utility function:

$$U = \log(W_\tau) + \sum_{t \leq \tau} z_t u^j \tag{3}$$

The first term in Equation (3) is the utility the investor derives from consuming her final wealth W_τ. The second term is the sum of flow utilities that the investor derives from enrolment, independently of her consumption. We assume that the flow utility is a constant u^j for each period of enrolment. Values of $u^j > 0$ can be generated through useful auxiliary services provided by the robo-advisor (tax strategies, for example), or, as we describe further below, through a "warm glow" of enrolment in the service. Conversely, negative parameter values $u^j < 0$ can capture disutility from enrolment in robo-advice, such as the psychological effects of algorithm aversion (e.g., Dietvorst, Simmons, and Massey (2015)).

Notice that the flow utility, u^j, can further depend on the type j of the investor's human advisor. This specification allows for several interpretations. For example, suppose that $u^1 > u^0$. On one hand, this could reflect a situation in which a type 1 advisor uses her greater interpersonal skill to alleviate algorithm aversion or distrust of automated solutions. On the other hand, a type 1 advisor could generate greater trust, or anticipation of higher risk-adjusted returns, in the sense of Gennaioli, Shleifer, and Vishny (2015), so that investors behaviorally enjoy a greater utility from risky investment

[15] This assumption eases exposition but does not materially affect our results, as long as one assumes that an investor who has quit cannot observe advised returns, or learn from them as described below, while she is not enrolled. Under this assumption, re-enrolment would typically be a dominated strategy in our model, because investors who have quit tend to have pessimistic beliefs about the quality of robo-advice and, in the absence of further learning, are unlikely to re-enrol.

when dealing with a type 1 advisor. The important common feature is that u^j, as specified in Equation (3), captures benefits from enrolment that are independent of wealth accumulation.

The investor's wealth prior to the end of her investment horizon evolves according to the following law of motion:

$$W_{t+1} = \left[z_t R_{t+1} + (1 - z_t) R^0_{t+1}\right] * W_t \tag{4}$$

If she remains enrolled at date t, i.e. $z_t = 1$, the investor earns a random gross return R_{t+1} on her wealth between periods t and $t+1$. If she quits, i.e. $z_t = 0$, she earns a return R^0_{t+1}, which represents her outside investment option.

Notice that this model focuses on the extensive margin of participation in robo-advice: In every period the investor either delegates her entire wealth to the robo-advisor, or invests her entire wealth in her outside option. The model can be generalized in a standard way to allow the investor to choose a fraction of wealth allocated to the robo-advisor. However, it is difficult to structurally estimate such a model in our empirical setting because we do not observe investors' total wealth outside of our data provider.

The returns R_{t+1} generated by robo-advice have a log-normal distribution. We use lower-case variables to denote natural logs, with $r_{t+1} \equiv \log(R_{t+1})$ and $r^0_{t+1} \equiv \log(R^0_{t+1})$. We then assume that excess log returns, denoted by r^e_{t+1}, satisfy the following relation:

$$r^e_{t+1} \equiv r_{t+1} - r^0_{t+1} = \theta + \varepsilon_{t+1}, \tag{5}$$

In Equation (5), the constant θ denotes the constant expected excess log return generated by robo-advice relative to the investor's outside option. The shock $\varepsilon_t \sim N(0, \sigma^2_\varepsilon)$ is drawn from a normal distribution, independently and identically across time periods, where the parameter σ_ε denotes the volatility of returns conditional on θ.

We interpret θ, which plays a key role in our analysis, as a measure of the *quality* of the portfolio choice algorithm underlying robo-advice. Indeed, in this environment, θ is the average log return that can be gained by enrolment, and therefore corresponds exactly to the pecuniary value of enrolment for an investor with log utility. As a broader interpretation, one can think of θ as a statistic summa-

rizing different dimensions of quality, for example, the robo-advisor's ability to correctly measure the investor's preferences and goals.

We assume that the data from which we estimate the model is generated by a robo-advisor with a fixed, true value of θ, regardless of the assignment j of human advisor types, and that this assignment is also independent of the realization of any shocks ε_t. Thus, consistently with the argument we made in the empirical analysis above, human advisors do not influence the return realizations of customer portfolios.

To close the model, we further assume that the investor is uncertain about the true quality θ of robo-advice. Indeed, she begins her life with the prior belief that quality has a normal distribution:

$$\theta \sim \mathcal{N}(\mu_0^j, 1/\tau_0^j) \tag{6}$$

Thus, the investor initially believes that the average excess return generated by the robo-advisor is μ_0^j, and the parameter τ_0^j denotes the precision of her belief. For instance, a confident belief that θ is large would be represented by large values for both μ_0^j and τ_0^j. Recall that θ is the mean of the *excess* log return that the robo-advisor generates relative to the investor's outside option. In principle, one could therefore interpret a high value of μ_0^j as capturing either an investor who is optimistic about robo-advice, or an investor who has a pessimistic expectation of her outside investment opportunities or skills. However, since we are interested in investors' empirical behavior while using robo-advice, we do not pursue the latter interpretation further. Therefore, we assume in the rest of this section that the distribution of the outside return r_{t+1}^0 does not depend on θ, and normalize it to an exogenous, constant value $r_{t+1}^0 = \bar{r}^0$.

We allow the parameters determining the investor's initial belief in Equation (6) to depend on the type j of the investor's human advisor. Once again, this specification permits both rational and behavioral interpretations. For a rational example, suppose that all investors begin life with a common belief about quality θ, and are shown a noisy piece of evidence by their human advisor, before forming the prior defined in Equation (6) using Bayes' rule. Suppose further that type 1 human advisors have access to more precise/less noisy evidence than type 0 advisors. In this example, a high value of μ_0^j corresponds to a favorable realization of the evidence, while a high value of τ_0^j corresponds to

an assessment that the evidence is precise. Rational expectations therefore lead to $\tau_0^1 > \tau_0^0$ in this scenario. For an alternative, behavioral example, one can simply interpret μ_0^j and τ_0^j in Equation (6) as the result of coaching, handholding, or persuasion by the human advisor.

Conditional on given values of μ_0^j and τ_0^j, the rational and behavioral interpretation yield observationally equivalent investor behavior. We therefore do not attempt to infer the source of prior beliefs from the data. However, conditional on prior beliefs, we assume that investors are fully rational, forward looking optimizers. Thus, the investor updates her beliefs using Bayes' rule based on realized excess returns, and correctly solve their dynamic optimization problem, which we outline below.

Bayes' rule in our model implies tractable updating equations. In particular, the log-normal distribution of returns in Equation (5) implies that the observation of each excess log return r_{t+1}^e gives the investor an unbiased Gaussian signal of the quality parameter θ. Hence, the investor's posterior beliefs remain Gaussian over time. We define these beliefs by $m_t^j = \mathbb{E}[\theta|R_1, \ldots, R_t]$, which denotes the investor's posterior expectation of excess returns based on information available at t, and by $\tau_t^j = 1/\mathbb{V}ar[\theta|R_1, \ldots, R_t]$, which denotes the precision of her posterior belief. We include the superscript j in these definitions because prior beliefs depend on the human advisor's type j, so that posterior beliefs may also depend on j.

The investor's beliefs evolve according to the following filtering equations:

$$m_{t+1}^j = m_t^j + \frac{\tau_\varepsilon}{\tau_t^j + \tau_\varepsilon} * (r_{t+1}^e - m_t^j) \tag{7}$$

$$\tau_{t+1}^j = \tau_t^j + \tau_\varepsilon \tag{8}$$

where $\tau_\varepsilon = 1/\sigma_\varepsilon^2$ denotes the inverse variance or precision of returns as a signal of quality θ. These two equations correspond to a standard Kalman filter. First, in Equation (7) the investor revises her expectation m_t of quality upwards whenever the newly observed realized return r_{t+1}^e is greater than her previous expectation. The weight placed on the new observation is given by the Kalman gain parameter $\frac{\tau_\varepsilon}{\tau_t^j + \tau_\varepsilon}$. This weight is greater when returns are a more precise signal of quality, i.e., when τ_ε is large. It is lower, however, when the investor's current belief is precise (we discuss this futher below), i.e., when τ_t^j is large. Second, in Equation (8), the investor increases the precision of her

beliefs linearly over time, with each observed return adding τ_ε units of precision.

An interesting implication is that confident prior beliefs about quality, in the sense that τ_0^j is large, imply a smaller response to realized returns in every subsequent period.[16] For example, if human advisors of type 1 instill more initial confidence, with $\tau_0^1 > \tau_0^0$, then the model predicts that investors matched with a type 1 advisors will be less responsive to returns in each subsequent period. In our structural estimation exercise below, we will use this feature to aid empirical identification of the prior precision parameter τ_0^j for different types of human advisor. Before turning to the empirical application, we now characterize optimal investor behavior.

3.2 Optimal Investor Behavior

The investor in our model solves an optimal stopping problem, deciding whether to stay enrolled in robo-advice or to quit in each period. At each date, the continuation/quitting decision can be made contingent on the history of returns observed so far. The investor chooses her strategy in order to maximize her expected lifetime utility.

In Appendix OA.1, we characterize the solution to this problem in recursive form. At any date t, the state variables in the investor's problem are her current wealth $w = w_t$, her current mean belief $m = m_t^j$ about the quality of robo-advice, and the precision $\tau = \tau_t^j$ of these beliefs. The evolution of these state variables is governed by investors' budget constraint in (4) and the belief updates in Equations (7) and (8). Consider an investor who is enrolled in robo-advice, and has been matched to a human advisor of type j. We define $V^j(w, m, \tau)$ as the investor's maximized value as a function of state variables, which captures her expected lifetime utility. In the appendix, we show that the value function takes the following shape:

$$V^j(w, m, \tau) = w + \frac{\bar{r}^0}{\rho} + F^j(m, \tau) \tag{9}$$

[16] This is because the difference equation in (8), when solved forward, yields $\tau_t = \tau_0^j + t * \tau_\varepsilon$. Thus, the Kalman gain parameter in Equation (7) is equal to $\frac{\tau_\varepsilon}{\tau_0^j + t * \tau_\varepsilon}$, which is decreasing in τ_0^j.

where $F^j(m,\tau)$ is a function defined recursively by

$$F^j(m,\tau) = \max\left\{0, \underbrace{m + u^j + (1-\rho)\mathbb{E}[F^j(m',\tau')|m,\tau]}_{\text{continuation value } \equiv C^j(m,\tau)}\right\} \tag{10}$$

Equations (9) and (10) have an intuitive interpretation. First, as a thought experiment, consider an investor with log wealth w who quits robo-advice and invests only in her outside option until the end of her horizon. The outside option yields the expected log return $\bar{r}^0$ per period and the investor has an average of $\frac{1}{\rho}$ periods left to invest.[17] Hence, the expected log final wealth of the investor who quits, which is also equal to her expected utility, is $w + \frac{\bar{r}^0}{\rho}$. This value is reflected in the first two terms in Equation (9).

Second, the function $F^j(m,\tau)$ defines the value of the option to remain enrolled in robo-advice, which is defined recursively in the Bellman equation (10). The option value is equal to the maximum of either zero (because the investor can always quit), or the continuation value $C^j(m,\tau)$ she can earn by remaining enrolled. The continuation value, in turn, has two components. First, by enrolling for one more period, the investor earns an expected excess return of m and enjoys flow utility u^j. Second, with probability $1-\rho$ she continues investing next period, and enjoys the expected future option value $\mathbb{E}[F^j(m',\tau')|m,\tau]$ of enrolment, where we use primes to denote next period's values of state variables. The continuation value in Equation (10) is the sum of the (expected) values of these two components.[18]

For a rational investor, it is optimal to quit whenever the continuation value $C^j(m,\tau)$ becomes negative. This option value does not have a closed-form, and we describe our numerical solution method in Appendix OA.2. However, the option value has intuitive properties, which we illustrate in Figure OA.VIII. The figure plots $m+u^j$, which is the per-period expected total payoff from enrolment,

[17] Because the probability that the investment horizon ends is ρ each period, the expected waiting time is $\frac{1}{\rho}$.

[18] The option value in Equation (10) obeys the terminal condition $\lim_{\tau\to\infty} F^j(m,\tau) = \frac{\max\{0,m+u^j\}}{\rho}$. This condition, derived in the appendix, follows by considering a rational investor who is certain about the quality of robo advice (i.e., the precision $\tau \to \infty$). Suppose this investor has a belief m about quality. She also knows that she will have $m' > 0$ in future periods, because she places zero weight on new information in Equation (7). If her (certain) estimate of quality satisfies $m + u^j < 0$, it is optimal to quit immediately and the option value is zero. If $m + u^j > 0$, then it is optimal to remain enrolled until the end of her horizon, which gives a total expected payoff (return plus flow utility) of $m + u^j$ per period, for an expected $\frac{1}{\rho}$ periods, so that the option value is $\frac{m+u^j}{\rho}$.

on the vertical axis, and the precision τ of investors' beliefs on the horizontal axis. The precision of beliefs on the horizontal axis is shown in units of the precision of returns, for example, $\frac{\tau}{\tau_\varepsilon} = 120$ denotes the point at which the investor's beliefs are as precise as they would be after observing 10 years of return data. For the numerical solution illustrated in the figure, continuation value $C(m, \tau)$ turns negative, and a rational investor quits, whenever $m + u^j$ falls below the black line. For low levels of precision, the investor is relatively forgiving and remains enrolled even for negative expected payoffs. This reflects the standard intuition that uncertainty (low precision) increases the expected future value of an option, in this case the option to remain enrolled. For high levels of precision, where the option value is more limited, the investor quits as soon as the expected payoff falls to zero.

In the next section, we describe our approach to using model-implied quit decisions to empirically identify the model's parameters.

3.3 Structural Estimation

In order to facilitate the structural estimation of the model, we begin by directly calibrating a subset of the model's parameters, which are summarized in Table 3 Panel A. We interpret one period in the model as one month. First, we assume that the data is generated by a robo-advisor with constant quality θ. We set $\theta = 0.0021$ to match the average difference between monthly log returns experienced by advised clients in our sample, and returns experienced by self-directed investors. Notice that this value reflects the better empirical performance among advised investors in our data, which we discussed in the context of Figure 1 above. Second, we set the volatility of returns $\sigma_\varepsilon = 0.00134$, which matches the monthly standard deviation of advised returns in the data. Third, based on the average age of investors in our sample, we set $\rho = 1/180$, implying an average investment horizon of 15 years (180 months). Finally, we impose that investors have unbiased expectations about θ ex ante, setting their prior expectations to $\mu^j = \theta$. We discuss the rationale behind the latter choice in more detail in the next section.

We estimate the remaining parameters from empirically observed investor behavior. Figure 8 shows Kaplan-Meier estimators of survival rates across the first 15 months of enrolment, using the same sample of investors as the estimates in Section 2.[19] We separately calculate these survival

[19]The estimates of conditional survival rates become noisy beyond 15 months of enrolment in our sample, due to

rates for low-retention human advisors (which we map to type 0 in the model), in the left-hand plot, and high-retention human advisors (mapped to type 1 in the model), in the right-hand plot. For each advisor type, we further condition survival probabilities on whether the investor experienced a negative (low) or positive (high) average holding return since enrolment. The dashed line in each plot is the point estimate in our sample, and the solid lines show 95% bootstrapped confidence intervals, which we describe in more detail below. Consistent with our earlier empirical results, the conditional survival curves show that type 1 human advisors achieve higher retention rates, and that the spread between retention rates after high and low returns is more pronounced for type 0 advisors.

We estimate model parameters using these survival curves as targets, deploying a minimum distance estimator (MDE).[20] More specifically, let $\hat{\pi}_{j,b,t}$ denote the estimated survival probability in Figure 8 for advisor types $j = 0, 1$, return bins $b =$ High, Low, and time period $t = 1, \ldots, 15$, for a total of 60 empirical moments. Let $g_{j,b,t}(\alpha)$ be the equivalent survival probabilities implied by the model, where α is the vector of model parameters to be estimated, which we define in detail below. We estimate the model parameters by minimizing the sum of squared deviations between empirical and model-implied survival probabilities. Thus, our estimates $\hat{\alpha}$ are given by:

$$\hat{\alpha} = \arg\max_{\alpha} \sum_{j} \sum_{b} \sum_{t} (\hat{\pi}_{j,b,t} - g_{j,b,t}(\alpha))^2 \tag{11}$$

Since we minimize an unweighted sum, we are essentially imposing an identity weighting matrix on the minimum distance estimator. Appendices OA.2 and OA.3 describe the relevant numerical procedures in detail. To briefly summarize, for any constellation α of parameters, we numerically solve the model using the Bellman equation (10) on a discretized grid, and calculate the model-implied survival probabilities $g_{j,b,t}(\alpha)$ using Kolmogorov forward equations. This procedure means that we are not employing a simulation-based or SMM estimator, thus eliminating an additional source of potential imprecision. We iterate over parameter constellations and use numerical optimization

censoring combined with the infrequent observation of negative returns. We therefore report unconditional survival estimates in empirical restuls in Figure 3, for example, but refrain from using their conditional equivalents as target moments for structural estimation.

[20] Since the Kaplan-Meier estimators that we use as empirical targets are not raw moments of the data, we are conducting minimum distance estimation, as opposed to a method of moments. In the terminology of MDE (e.g., Cameron and Trivedi (2005), Chapter 6.7), the survival probabilities are *reduced-form parameters* that we use to infer the deep parameters of the model.

to solve the estimation problem in (11). Before reporting our estimates, we discuss the sources of structural identification of our estimated parameters.

3.4 Identification of Model Parameters

In this section, we discuss the logic behind the identification of the model parameters. Appendix OA.4 contains further analytical and empirical results supporting our intuitive arguments. The estimated model parameters are investors' prior beliefs about quality θ, parametrized by mean μ^j and precision τ^j, as well as investors' flow utility u^j from enrolment. Each of these is defined for both types $j \in \{0, 1\}$ of human advisor. In addition, we estimate a fitting parameter denoted η, which we introduce in order to generate cross-sectional heterogeneity in investor quit decisions. In particular, as is common in structural applications with discrete choice, we assume that quit decisions in the model are probabilistic, and driven by the following rule:

$$\Pr\left[\text{quit}|m, \tau, \text{enrolled}\right] = \frac{1}{1 + \exp\left\{-\frac{C^j(m,\tau)}{\eta}\right\}} \tag{12}$$

This assumption imposes a logistic choice rule whereby the investor quits more frequently when the continuation value from enrolment is low. One possible foundation for Equation (12) is that at each date t, the investor in the model acts as if her flow utility from enrolment over the next period is $u^j + \eta \cdot \xi_t$, where ξ_t stands in for the variety of unmodeled reasons for remaining enrolled, which could include, for example, liquidity shocks or implementation error. The fitting parameter $\eta > 0$ measures the noisiness of quit decisions. For example, if $\eta = 0$, the investor quits with probability 1 if $C^j(m, \tau) < 0$ and with probability 0 otherwise. By contrast, if $\eta \to \infty$, then investors quit with 50% probability regardless of continuation values.[21]

The survival curves in Figure 8 allow us to base structural identification both on the *level* of survival rates for different advisor types, and on the *sensitivity* of survival rates to realized returns. Four intuitive observations—made rigorous in the appendix—explain the sources of structural identification. First, for every investor type j, the level of survival rates identifies the flow utility parameter u^j, since

[21] While we estimate η below to maximize the fit of our model to the data, Appendix OA.3.1 shows that our qualitative results are robust to different choices of η within a reasonable range.

lower utility from enrolment translates to higher quit rates/lower survival probabilities. Second, for every type j, the sensitivity of survival rates to returns identifies the parameter τ_0^j measuring the initial precision of beliefs. Indeed, recall from Equations (7) and (8) that investors place a greater weight on realized returns when updating their beliefs if the prior precision is low. Therefore, an investor with lower prior precision is more likely to be driven below the quit threshold conditional on experiencing low realized returns.

Third, the sensitivity of survival rates to returns is also used to identify the fitting parameter η in Equation (12), since higher values of η mean that quitting decisions will be driven by pure noise and, conversely, will be less dependent on realized performance. We are able to identify this parameter separately from the precision parameters τ_0^j for two reasons. On the one hand, we force η to be independent of human advisor type, which conservatively reduces the model's degrees of freedom. On the other hand, while the precision τ_0^j drives the sensitivity of belief updates to returns, it mostly does so in early periods after enrolment, because the weight placed on returns in later periods (when the investor has already learned from a longer time series) is muted. Therefore, heuristically, the sensitivity of survival rates to returns in early periods identifies τ_0^j, while the sensitivity in later periods separately identifies η.

Finally, one can apply similar reasoning to argue that the prior mean μ^j is identified primarily by the level of survival rates in early periods. In early periods, investors have not observed many returns and still place a greater weight on their prior beliefs, so that a change in μ^j has a substantial effect on quit decisions. Thus, μ^j can be identified separately from the flow utility parameter u^j for each type of human advisor, which governs the level of survival rates in late periods. However, in order to reduce our reliance on the distinction between early and late periods, we discipline the model by assuming that $\mu^0 = \mu^1 = \theta$ for each advisor type. Economically, this assumption reflects a situation in which all investors have access to an unbiased point estimate that reflects the historical mean of advised returns in our sample. The precision of beliefs τ_0^j then reflects the amount of initial confidence in the point estimate, which is allowed to differ across human advisor types.

To summarize, we use the $4 \times 15 = 60$ estimated survival probabilities in Figure 8 in order to identify five model parameters, which we collect in the vector $\alpha = (\eta, \tau^1, u^1, \tau^2, u^2)$. The next section

discusses the results of estimation.

3.5 Estimation Results

In Table 3 Panel B, we report our estimates of model parameters. For each parameter, the table reports its estimated value, standard error and a t-statistic. The standard errors, which we describe in detail in Appendix OA.3, are based on a bootstrapped estimate of the variance-covariance matrix of the empirically estimated survival probabilities $\hat{\pi}_{j,b,t}$, and are clustered by time period in order to control for common shocks that affect all investors' decisions in any given month.

The first row shows our estimate of the common fitting/noise parameter $\eta = 0.2010$, which has no direct economic interpretation but is precisely estimated. The following rows show the advisor type-specific estimates of the standard-deviation of beliefs σ^j and flow utilities u^j. The estimates of the volatilities of beliefs (in what follows, we discuss the reciprocals of belief variances, i.e., precisions, for ease of interpretation) are reasonably tightly estimated with t-statistics of 11 and 5 respectively, but remain noisier than the estimates of flow utilities, with t-statistics of 34 and 26. Our estimates therefore reveal two dimensions of the structural differences between type 0 and type 1 human advisors.

3.5.1 The Structural Effects of Human Advisors

First, type 1 advisors (i.e., high-retention types) induce substantially higher precision in investors' initial beliefs about the quality of robo-advice: $\tau^1 = 1/(\sigma_0^1)^2$ is around 66% higher than $\tau^0 = 1/(\sigma_0^0)^2$. In order to interpret the economic significance of this difference, one can express precisions in units of return observations, using the calibrated precision of returns τ_ε in Panel A. The precision induced by a type 1 advisor is equivalent to observing a time series of $1289396/5487 \simeq 235$ monthly returns or roughly 20 years of performance data, at the start of enrolment.[22] By contrast, the precision induced by a type 0 advisor is equivalent to roughly 12 years of performance data. This result suggests that high-type human advisors add value by increasing their clients' confidence in product quality.

Second, type 1 advisors deliver greater flow utility to investors, since u^1 is around 6% higher than u^0. For an economic interpretation, the difference in flow utilities yields a lifetime benefit that is

[22] More rigorously, a quasi-Bayesian investor who begins life with a diffuse prior about θ, observes 235 monthly returns, and updates from each return using Equations (7) and (8), will end up with a posterior belief with precision close to τ_0^1.

equivalent to a portfolio return of 0.3% return (30bps) per annum. The baseline estimate u^0 is still positive, which has two possible interpretations. On the one hand, the positive value could reflect the benefit of other services that our studied robo-advisor provides, such as income or estate tax planning strategies, or other financial management. On the other hand, recall that we have forced investors' prior expectation μ_0^j to be equal to the historical mean of returns in the data. Another interpretation of the positive estimated flow utilities is therefore that investors act as if they expect higher risk-adjusted returns than implied by historical performance (Gennaioli, Shleifer, and Vishny, 2015).

3.5.2 Implications for Communication and Logins

These estimates are also interesting against the backdrop of our empirical analysis of client-advisor interactions, which showed that high-retention advisors spend significantly more time (particularly in normal times) initiating and conducting communication with clients. The structural model further indicates that this communication may be associated with meaningful effects on client confidence and overall utility while using robo advice. Moreover, we showed that the clients of high-retention advisors spent less of their own time logging into the robo-advisor's web platform. Once again, this behavior is consistent with investors who are more confident in their beliefs about the service and therefore do not feel a frequent need to check their accounts.

3.6 Value of Complementarity and Counterfactuals

In this section, we provide further economic context for our structural results by quantifying the surplus obtained by investors, as well as by the advisory firm, in different advice scenarios. We also consider counterfactual scenarios that further illustrate the sources of economic gains.

Table 4 collects the associated results. In the first two columns, for each type of human advisor, we calculate the surplus from enrolment for the investor and the firm, conditional on the true value of quality θ. We explain the calculation of investor and firm surplus rigorously in Appendix OA.1. For the investor, we calculate surplus as the difference between the investor's expected lifetime utility if she is initially enrolled in robo advice and then follows the quit rule in Equation (12), and her

expected lifetime utility if she earns only self-directed returns until the end of her investment horizon. For example, for an investor with a type 0 advisor, expected surplus is given by 0.6068 units of utility, while the surplus achieved by an investor with a type 1 advisor is around 15% higher at 0.6981. In money-metric terms, the investor with a type 0 advisor would obtain the same utility if i) she is reassigned to a type 1 advisor, or if ii) she receives a lump-sum transfer equal to 9.6% of her initial wealth.[23]

For the firm, we calculate surplus as the expected value of fees earned from investor enrolment, given that the investor follows the quit rule in Equation (12). We set fees to 0.3% of assets under management per annum, and assume that the firm discounts future cash flows from these fees at a rate of 0.5% per annum, since the Fed Funds rate ranged between zero and 1% over our sample period.[24] We report increases in firm surplus, defined as the expected net present value of future fees, per dollar of assets under management. The difference between the firm's surplus for a type 0 versus type 1 human advisor, which results from the lower quit rates induced by type 1 advisors, is around 0.16% (26bps) per dollar of investor assets. While this figure is modest in per-dollar terms, it can potentially add up to a significant effect given the scale of major financial advisory firms.

In the third and fourth column of Table 4, we report the model-implied investor and firm surplus in two counterfactual scenarios. The "Type 0 + more information" scenario keeps the flow utility provided by a type 0 advisor but increases the precision of the investor's beliefs to τ_0^1, the value induced by a type 1 investor. The differences reported in the table are relative to the baseline type 0 advisor, and show only modest improvements, on the order of 0.4% of initial wealth for the investor and 1bp of assets under management for the firm. The "Type 0 + more utility" scenario keeps the precision of a type 0 advisor but raises the flow utility to u^1. This scenario results in a gain equivalent to 8.9% of initial wealth for the consumer, and 24bps for the firm.[25]

These counterfactual scenarios suggest that around 90% of consumer and firm gains from having a type 1 advisor can be generated by replicating the flow utility associated with type 1 advisors, without

[23]Given log utility, the lump sum transfer X of initial wealth that keeps the investor's expected utility unchanged relative to obtaining additional surplus S solves the equation $\log(W_0 + X) = \log(W_0) + S$, or $1 + \frac{X}{W_0} = \exp(S)$. For the surplus from being reassigned we use $S = 0.0961$.

[24]For ease of interpretation, we assume that the firm's cash-flows from fees are obtained with certainty.

[25]The gains in the two scenarios do not exactly add up to the overall difference between type 0 and type 1 advisors, which are reported in the first two columns, because of interactions between the model's parameters. For instance, the impact of a change in u^j on quit probabilities can depend on the baseline level of τ_0^j, and vice versa.

increasing investors' confidence in quality. Therefore, while the gain in confidence appears economically significant, it is less important in welfare terms. One interpretation of this result is that, to the extent that some investors experience algorithm aversion or distrust of automation, adroit advisors can help to alleviate it. Another is that investors derive substantial additional "warm glow" utility à la Gennaioli, Shleifer, and Vishny (2015) from interactions with their advisors. Both interpretations are interesting, and regardless of which one is correct, the magnitudes that we estimate suggest that firms would benefit greatly from further study and measurement of the underlying root causes of human advisor effectiveness.

4 Conclusion

We study the extent to which human experts are complementary to technological innovation using a simple model which we structurally map to empirical evidence. In the model, investors experience a per-period fixed cost or disutility arising from their use or interaction with the technology, as well as a learning channel, in which they refine their understanding of the performance of the technology over time and across states of the economy. In a unique dataset from a large US hybrid robo-advisor, we leverage random assignment of human advisors to clients and clients, study subsequent retention rates of clients in the service, and map these patterns in the data back to the parameters of the structural model.

The robo-advisor we study automatically manages the investment portfolio using a set of codified rules, while the human advisor interacts with investors to help them understand these automated investment decisions and how they help further the investors' financial goals, as well as advising on related servies such as tax strategies and debt management. A key feature of this setting is that the assignment of investors to advisors follows mechanical rules driven by workload balancing imperatives rather than any assessment of advisor type. This means that once the current client-load of a given advisor is accounted for, the assignment of new clients to this advisor is orthogonal to the historical client retention of the advisor (a useful proxy for advisor type).

This random assignment of clients to different types of advisors allows us to cleanly map our empirical estimates to the parameters of the model. We find that this measure of advisor type

predicts the future retention rate of clients that are assigned to them. We also find that high-retention advisors' clients are less likely to quit robo-advising during periods of market turbulence than lower-retention advisors' clients. Finally, we find that experienced clients, regardless of their advisor type, react less to market turbulence. These facts, when mapped back to the model, deliver the insight that humans are complementary to automated services in two main ways. For one, the estimates imply that high-quality human advisors help to reduce the variance in clients' prior beliefs about service quality, facilitating learning about the algorithm's ability to deliver returns. Second, the behavior of the attrition rates of experienced clients shows that human advisors can also significantly attenuate ongoing disutility from the automated portfolio management solution.

Our results reinforce the view that human social skills have high residual value in an environment of increasing automation, and make the underlying economics more transparent, separating the relevant effects into both learning and disutility components. Our counterfactual analysis suggests that, with judicious matching of advisors to clients, human advice can help to complement automation and the use of algorithms to facilitate the broader scaling of customized solutions in household finance and other domains.

References

Andries, M., M. Bonelli, and D. A. Sraer. 2024. Financial advisors and investors' bias. *Available at SSRN* .

Badarinza, C., J. Y. Campbell, and T. Ramadorai. 2016. International comparative household finance. *Annual Review of Economics* 8:111–44.

Balasubramaniam, V., J. Y. Campbell, T. Ramadorai, and B. Ranish. 2022. Who owns what? a factor model for direct stock holding. *Journal of Finance* .

Barber, B. M., and T. Odean. 2001. Boys will be boys: Gender, overconfidence, and common stock investment. *Quarterly Journal of Economics* 116:261–92.

Calvet, L. E., C. Célérier, P. Sodini, and B. Vallee. 2023. Can security design foster household risk-taking? *The Journal of Finance* 78:1917–66.

Cameron, C. A., and P. K. Trivedi. 2005. *Microeconometrics: Theory and applications*. Cambridge University Press.

Campbell, J. Y. 2006. Household finance. *The Journal of Finance* 61:1553–604.

Cao, S., W. Jiang, J. L. Wang, and B. Yang. 2021. From man vs. machine to man+ machine: The art and ai of stock analyses. *Columbia Business School Research Paper* .

Cao, S., W. Jiang, B. Yang, and A. L. Zhang. 2023. How to talk when a machine is listening: Corporate disclosure in the age of ai. *The Review of Financial Studies* 36:3603–42.

Collinson, R., J. E. Humphries, N. S. Mader, D. K. Reed, D. I. Tannenbaum, and W. Van Dijk. 2022. Eviction and poverty in american cities. *Quarterly Journal of Economics* forthcoming.

D'Acunto, F., N. Prabhala, and A. Rossi. 2019. The promises and pitfalls of robo-advising. *Review of Financial Studies* .

D'Acunto, F., and A. G. Rossi. 2022. Robo-advising. *Palgrave Macmillan Handbook of Technological Finance* .

Dietvorst, B. J., J. P. Simmons, and C. Massey. 2015. Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144:114–.

Egan, M., G. Matvos, and A. Seru. 2019. The market for financial adviser misconduct. *Journal of Political Economy* 127:233–95.

Gargano, A., and A. G. Rossi. 2017. Does it pay to pay attention? *Working Paper* .

Gennaioli, N., A. Shleifer, and R. Vishny. 2015. Money doctors. *The Journal of Finance* 70:91–114.

Gomes, F., M. Haliassos, and T. Ramadorai. 2021. Household finance. *Journal of Economic Literature* 59:919–1000.

Guiso, L., P. Sapienza, and L. Zingales. 2008. Trusting the stock market. *the Journal of Finance* 63:2557–600.

Guiso, L., and P. Sodini. 2013. Household finance: An emerging field. In *Handbook of the Economics of Finance*, vol. 2, 1397–532. Elsevier.

Linnainmaa, J. T., B. Melzer, A. Previtero, and S. Foerster. 2018a. Financial advisors and risk-taking. *Working Paper* .

———. 2018b. Financial advisors and risk-taking. *Unpublished manuscript* .

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

Politis, D. N., and J. P. Romano. 1994. The stationary bootstrap. *Journal of the American Statistical association* 89:1303–13.

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

Reher, M., and C. Sun. 2019. Automated financial management: Diversification and account size flexibility. *Journal of Investment Management* .

Reuter, J., and A. Schoar. 2024. Demand-side and supply-side constraints in the market for financial advice. *NBER Working Paper* .

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

Rossi, A. G., and S. Utkus. 2019. The needs and wants in financial advice: Human versus robo-advising. *Working Paper* .

Sicherman, N., G. Loewenstein, D. Seppi, and S. Utkus. 2016. Financial attention. *Review of Financial Studies* 29:863–97.

Tauchen, G. 1986. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics letters* 20:177–81.

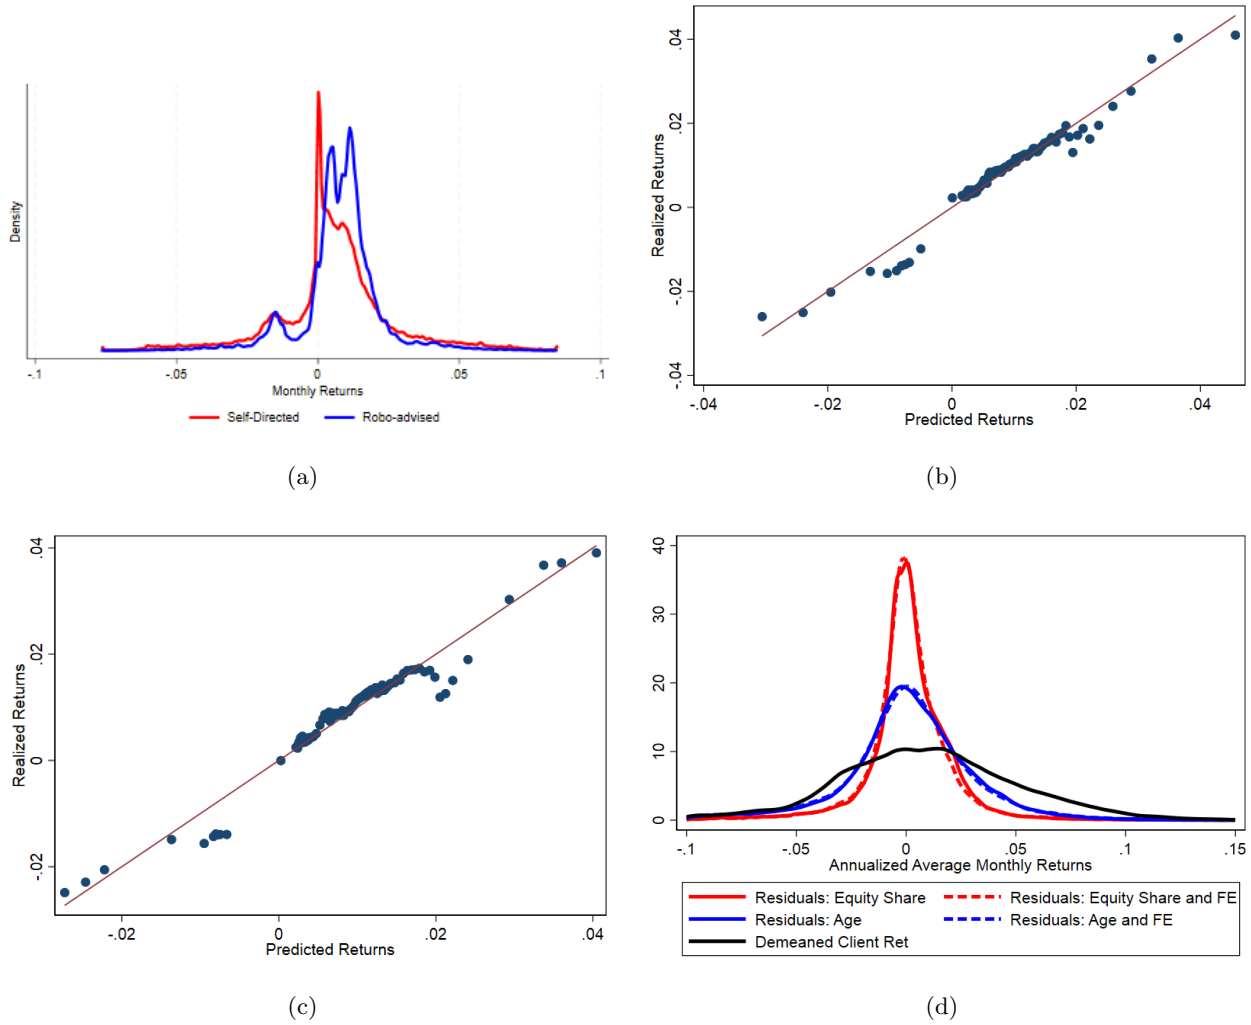
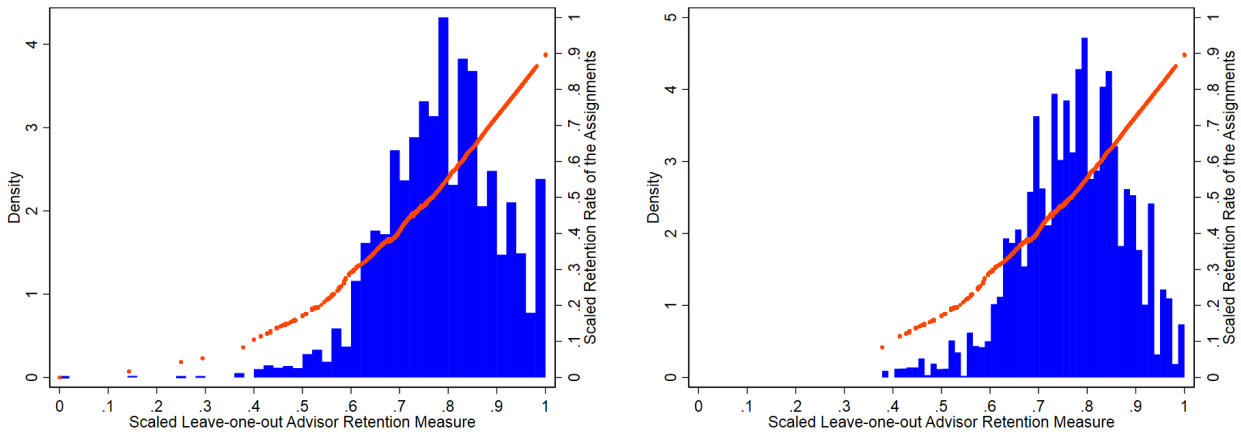
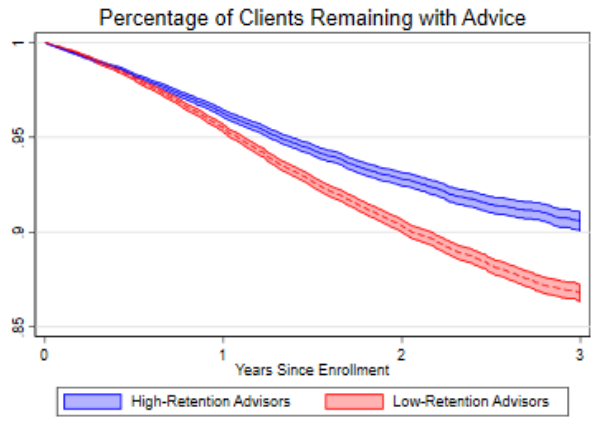


Figure 1: This figure reports results on investors' portfolio performance. We report the density of monthly returns across investors and periods for advised and non-advised individuals in subfigure (a). Subfigures (b) and (c) relate monthly client returns to predicted returns on the basis of investors' equity share and age, respectively. Subfigure (d) shows the extent to which investor returns are determined by their equity share, demographic characteristics, and their assigned advisor.

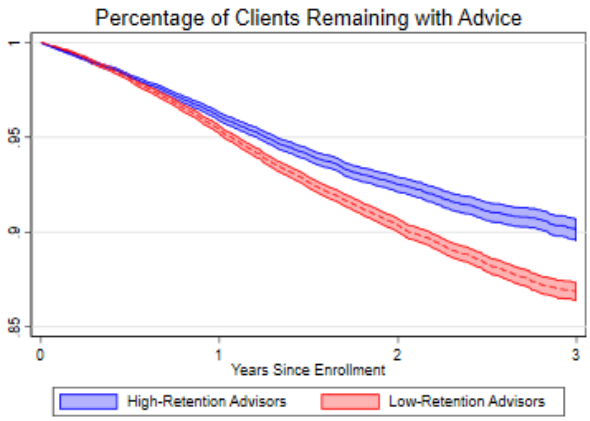


(a) (b)

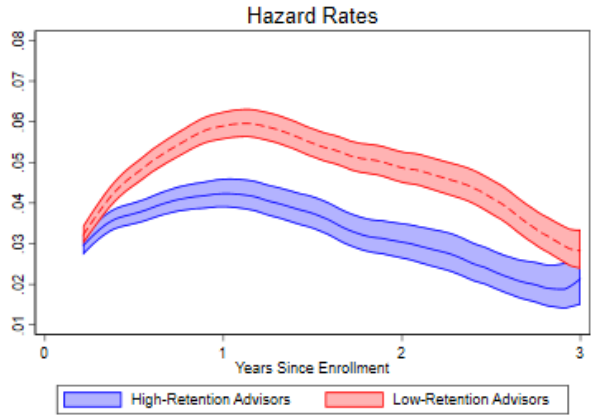
Figure 2: In subfigure (a), the histogram reports the percentage of clients retained by each advisor-client pairing, scaled in such a way that the advisor-client pairings with the highest retention are assigned a value of 100. Superimposed on the histogram, we report non-parametric estimates of the relation between the scaled leave-one-out retention measure of each advisor-client pair on the scaled retention rate of that specific client. Subfigure (b) repeats the exercise excluding the advisors that are in the bottom decile ranked by the number of advised clients.



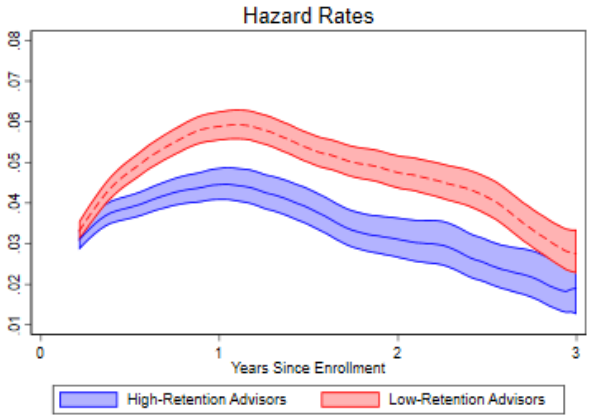
(a) All Advisors



(b) Excluding Lowest 10% of Advisors

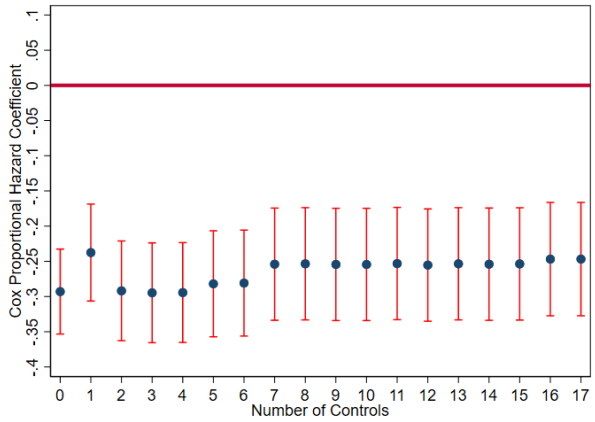


(c) All Advisors

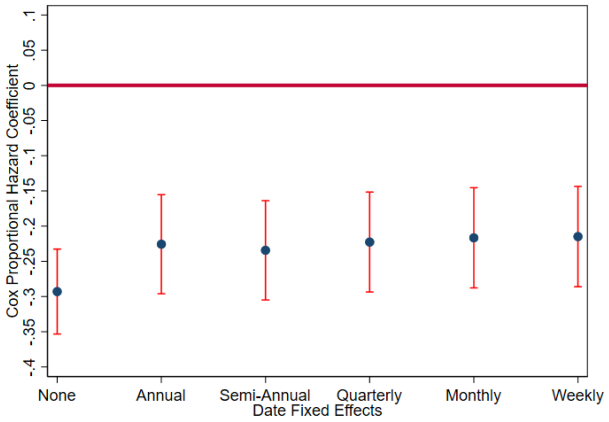


(d) Excluding Lowest 10% of Advisors

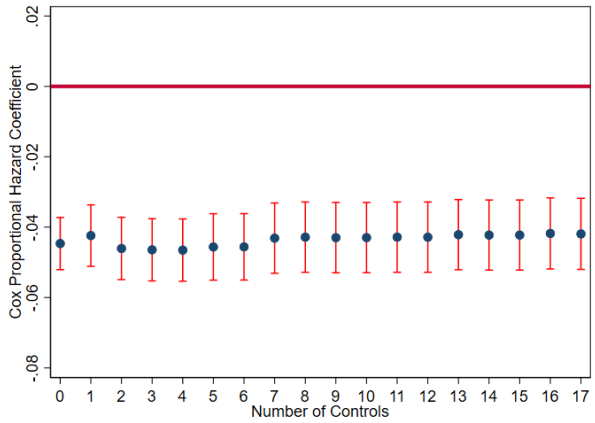
Figure 3: Subfigure (a) shows survival plots for investors assigned to advisors of different types. Subfigure (b) repeats the survival estimates excluding advisors in the bottom decile in terms of load. Subfigure (c) shows smooth hazards for investors assigned to advisors of different types. Subfigure (d) repeats the hazard estimates excluding advisors in the bottom decile in terms of load. Investors associated with high-retention advisors are in blue. Those with low-retention advisors are in red.



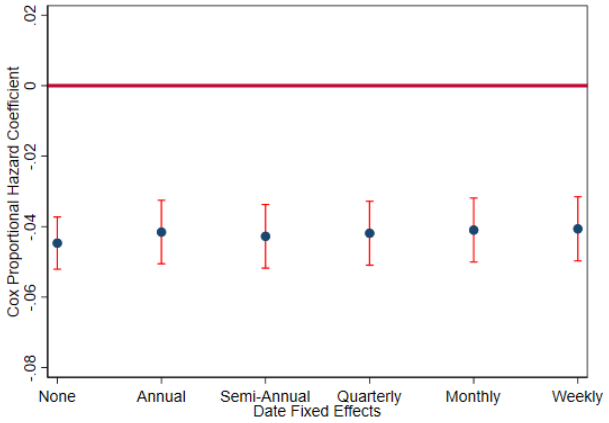
(a) Main Regressor: High-retention Dummy



(b) Main Regressor: High-retention Dummy



(c) Main Regressor: Advisor Retention



(d) Main Regressor: Advisor Retention

Figure 4: This figure reports estimates from Cox proportional-hazard models of the form: $h(t|x_j) = h_0(t) \cdot \exp(x_j\beta)$, estimated on the full set of advisors. Subfigure (a) presents coefficient estimates and 95% confidence intervals for the "high-retention" advisor dummy as client covariates from Table 1 are added sequentially in decreasing order of their absolute t-statistics—starting with the least balanced characteristics. Subfigure (b) re-estimates the analysis from (a) while progressively incorporating more granular time fixed effects, ranging from annual to weekly. Subfigures (c) and (d) re-estimate the analyses from (a) in (b), respectively, replacing the high-retention dummy with a continuous measure of advisor retention.

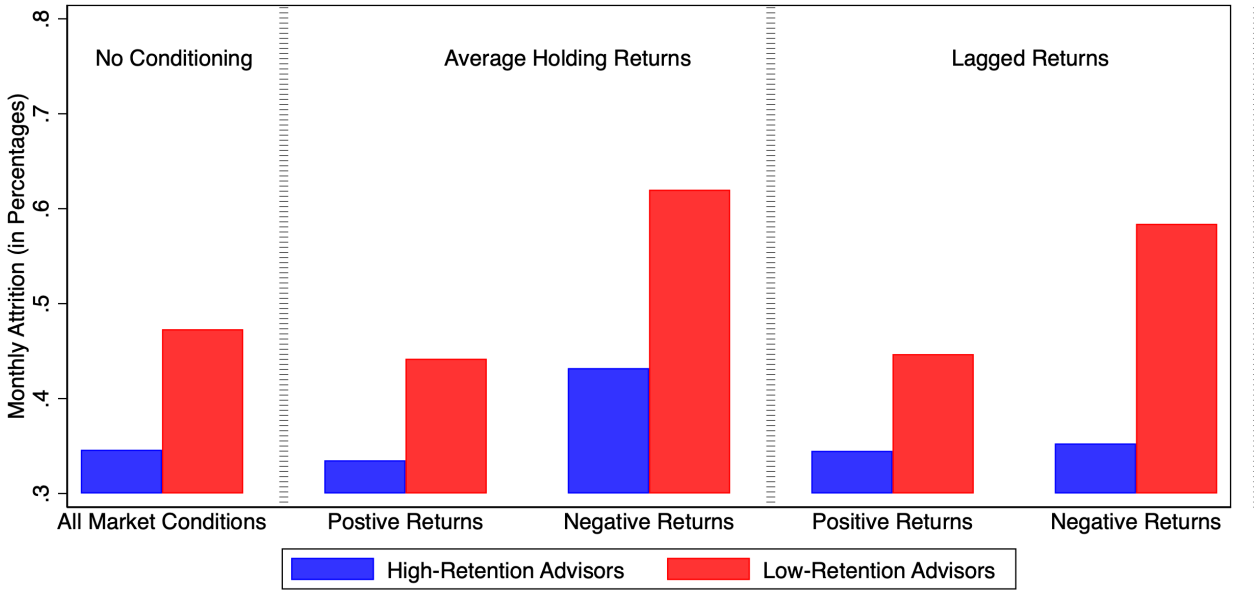


Figure 5: This figure displays the monthly attrition rates (in percentages) of investors assigned to different types of advisors. The first set of results does not condition on portfolio performance. The second and third sets of results condition on performance using lagged average portfolio returns since enrollment calculated from the month investors subscribe to robo-advice. The fourth and fifth sets of results condition on performance using one-period lagged portfolio returns.

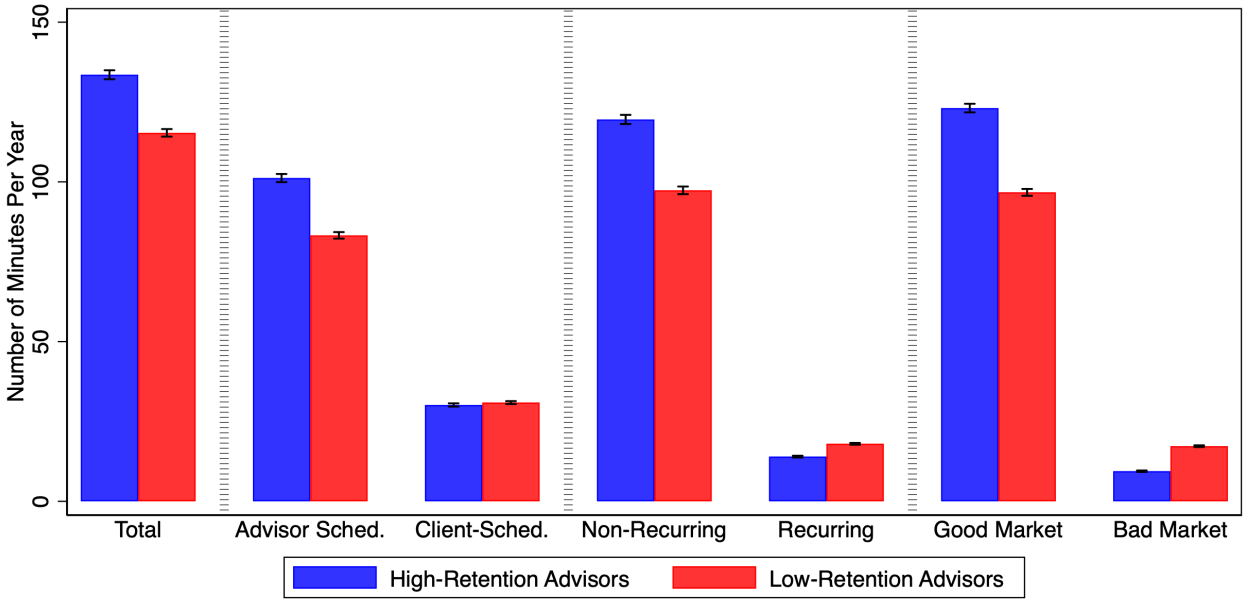


Figure 6: This figure reports the average number of minutes per year investors spend with their advisors, together with 95% confidence intervals. The first set of bars reports the results across all meetings. The second and third sets of bars compute the results separately for advisor- and client-scheduled meetings, respectively. The fourth and fifth sets of bars compute the results separately for non-recurring and recurring meetings, respectively. The last two sets of bars compute the results across good and bad market conditions.

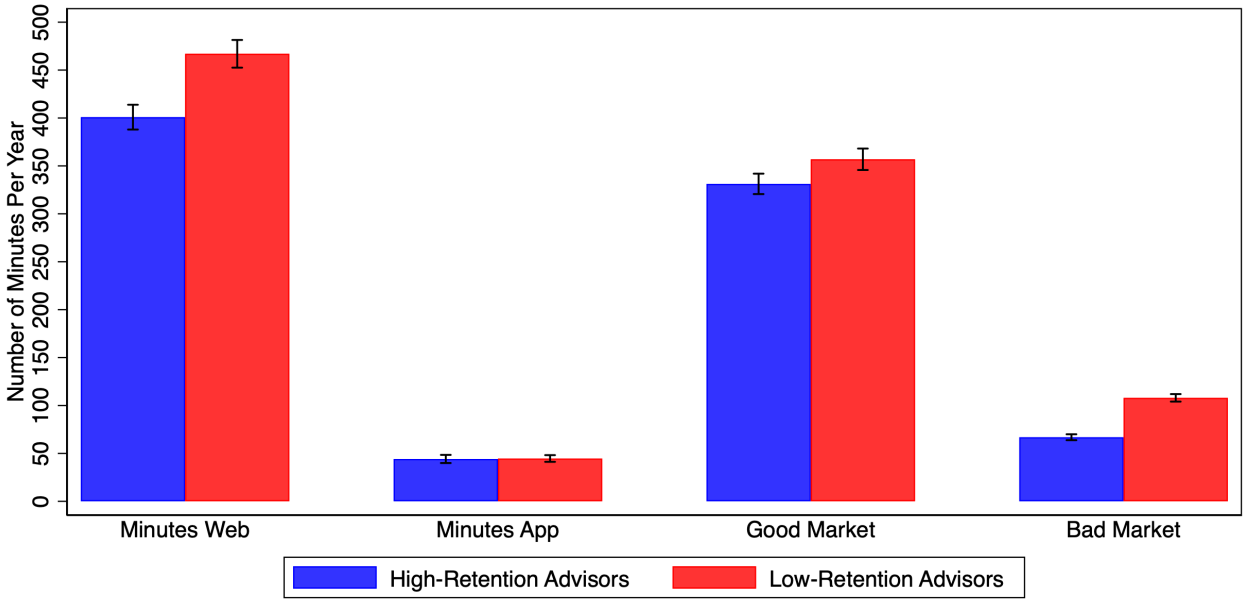


Figure 7: This figure reports the average number of minutes per year investors spend logged into their robo-investment portfolios, together with 95% confidence intervals. The first and second sets of bars compute the results separately time spent on the website and time spent on the app. The third and fourth sets of bars compute the results across good and bad market conditions.

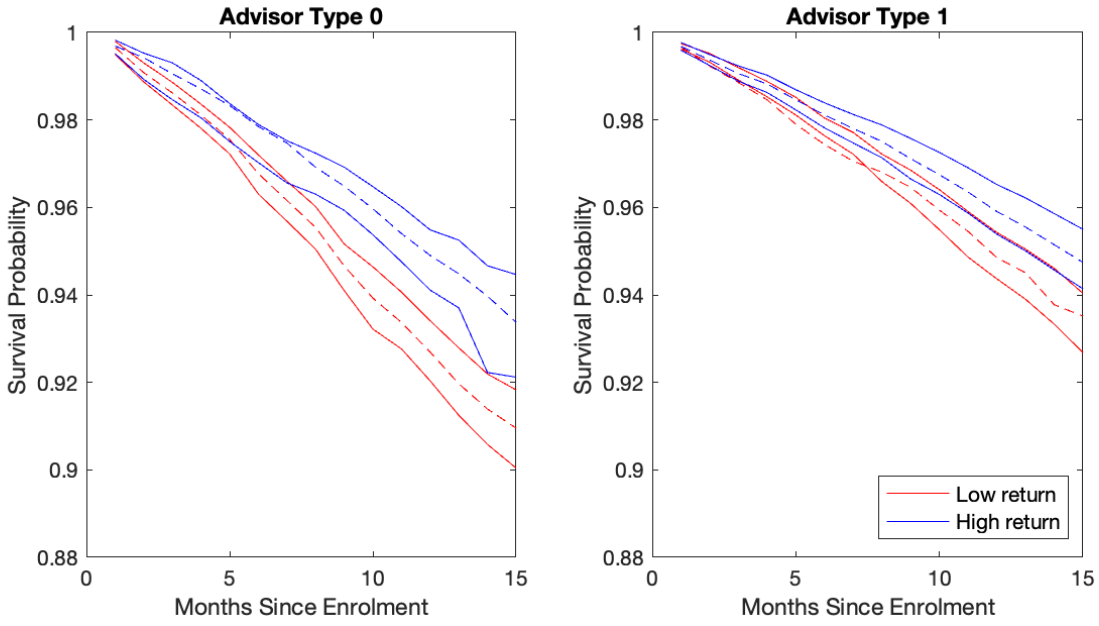


Figure 8: The left panel shows survival curves among investors assigned to a type 0 (low-retention) advisor, conditional on having experienced average returns below zero ("High") or above zero ("Low") since enrolment. The right panel shows the same survival curves for investors assigned to a type 1 (high-retention) advisor. The dashed line is the point estimate, and the solid lines define 95% bootstrapped confidence intervals, with clustering by time period. Since each curve is drawn up to 15 months since enrolment, these survival estimates define the $4 \times 15 = 60$ reduced-form parameters that we use for structural estimation.

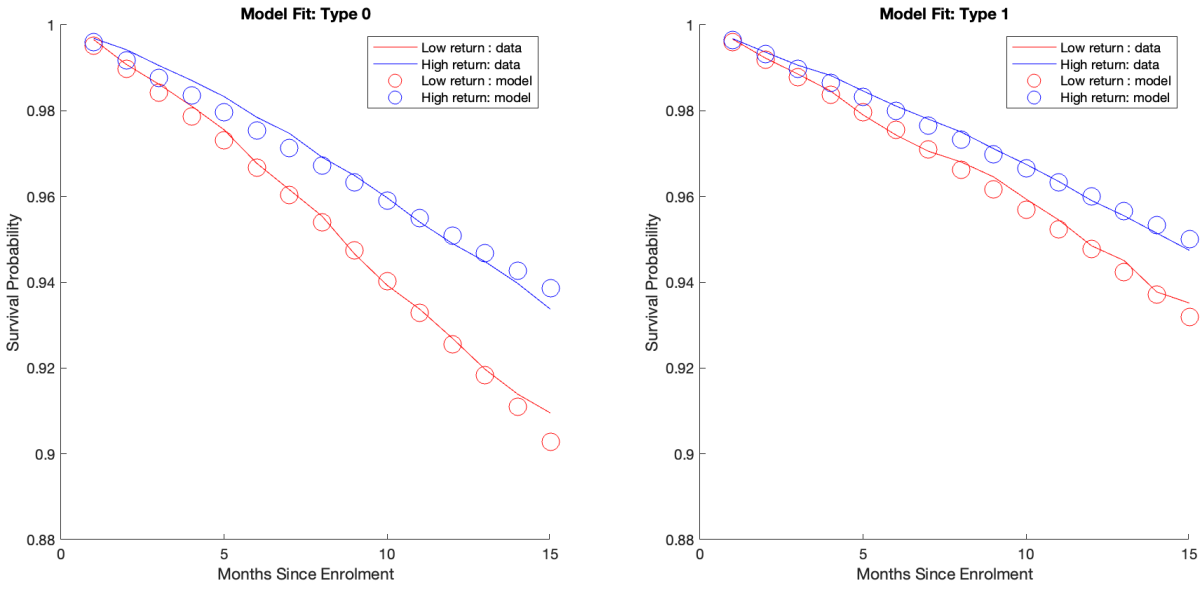


Figure 9: This figure shows the same empirical survival curves as Figure 8, i.e., conditional on advisor types 0 (low-retention) and 1 (high-retention) as well as High and Low returns. In each plot, the large circles show the associated model-implied survival rates, evaluated at the estimated parameter values.

Table 1. Covariates Balancing Across Clients to High- and Low-Retention Advisors

	High Retention		Low Retention		Diff		
	mean	N	mean	N	mean	t-stat	N
Age	64.540	24,514	65.657	23,511	1.117***	(3.33)	48,025
Male	0.577	25,739	0.618	24,085	0.041***	(4.76)	49,824
Tenure	14.666	25,739	15.635	24,085	0.969***	(3.03)	49,824
Wealth	946,754	25,739	993,861	24,085	47,107	(0.39)	49,824
NumAssets	10.717	25,739	11.438	24,085	0.721	(1.69)	49,824
PctAMProducts	0.853	25,706	0.850	24,062	-0.003	(-0.81)	49,768
PctMutualFunds	0.666	25,706	0.672	24,062	0.006	(0.61)	49,768
PctCash	0.234	25,706	0.226	24,062	-0.007	(-0.67)	49,768
PctETF	0.035	25,706	0.034	24,062	0.000	(-0.27)	49,768
PctStocks	0.046	25,706	0.047	24,062	0.001	(0.84)	49,768
PctBonds	0.002	25,706	0.002	24,062	0.000**	(2.65)	49,768
AcctIndex	0.436	25,739	0.438	24,084	0.002	(0.14)	49,823
MgtFee	0.147	23,877	0.147	22,931	0.001	(0.28)	46,808
ExpRatio	0.209	23,299	0.206	22,396	-0.003	(-0.18)	45,695
TurnRatio	0.328	22,918	0.343	21,787	0.016**	(2.16)	44,705
Ret. Pre-Robo	0.051	22,040	0.045	20,884	-0.005	(-0.98)	42,924
Adj. Ret. Pre-Robo	-0.007	22,040	-0.009	20,884	-0.002	(-1.64)	42,924

This table reports balancing results for demographic characteristics and portfolio allocation behavior for investors 1 month before signing up for advice. For each characteristic, in the first four columns we report the mean and the number of observations for high- and low-retention advisors. In the last three columns we report the difference in means, the associated t-statistic and the total number of observations.

Table 2. Human Advice and Investment Performance

	No Conditioning	Lagged Average Returns	Lagged One-Period Returns		
			All	Short Tenure	Long Tenure
Negative Return		0.178***	0.137***	0.157***	0.139**
		(4.01)	(2.98)	(2.99)	(2.23)
HighRetn_Advisor	-0.127***	-0.107***	-0.102***	-0.086***	-0.107***
	(-6.32)	(-5.34)	(-4.81)	(-2.67)	(-4.28)
Interaction		-0.081*	-0.129***	-0.120**	-0.192***
		(-1.88)	(-3.45)	(-2.22)	(-3.51)
Constant	0.473***	0.442***	0.447***	0.419***	0.466***
	(15.03)	(13.14)	(12.43)	(12.44)	(10.74)
Clustering	Date&User	Date&User	Date&User	Date&User	Date&User
R-square	0.00010	0.00017	0.00014	0.00017	0.00012
N	884,022	884,022	884,022	459,690	424,332

This table reports coefficient estimates of the following baseline regression estimated at the monthly frequency:

$$Dummy_quit_{i,t} = \alpha + \beta\ I_{\{Negative\ Return_{t-1}=1\}} + \gamma\ I_{\{HighRetn_Advisor_i=1\}} + \delta\ I_{\{Negative\ Return_{t-1}=1\}} \times I_{\{HighRetn_Advisor_i=1\}} + \epsilon_{i,t},$$

where $Dummy_quit_{i,t}$ is equal to 1 if investor i quits robo-advising in month t and 0 otherwise, $I_{\{Negative\ Return_{t-1}=1\}}$ is an indicator variable equal to 1 if the market return is negative in month $t-1$ and 0 otherwise, and $I_{\{HighRetn_Advisor_i=1\}}$ is equal to 1 if investor i is assigned to an advisor with client retention above the median and 0 otherwise. We multiply all the coefficient estimates by 100, so they are expressed in percentages. We use investors' lagged average returns since enrollment and investors' lagged one-period returns in columns (2) and (3-5), respectively. The sample of advised users is split into two groups on the basis of their tenure and report the results for short tenure in column (4) and long tenure in column (5). The standard errors are double-clustered by users and dates.

Table 3. Structural Model Parameters

Panel A. Calibrated Parameters

Parameter	Notation	Value	Target
Mean excess return	θ	0.0021	Avg. robo return - avg. self-directed return
Return precision	τ_ε	5487	St.dev. of robo return = $1/\sqrt{\tau_\varepsilon}$
Probability of horizon end	ρ	0.0056	15y (180m) average horizon = $1/\rho$
Prior mean	$\mu^0 = \mu^1$	0.0021	Mean Robo return - mean self-directed return

Panel B. Estimated Parameters

	Parameter	Notation	Value	Standard Errors	T-Statistic
Common	Noise/fitting parameter	η	0.2010	0.0046	44
Type 0 Advisor	Prior standard deviation	σ_0^0	0.00114	0.00010	11
	Flow utility	u^0	0.00396	0.00012	34
Type 1 Advisor	Prior standard deviation	σ_0^1	0.00088	0.00017	5
	Flow utility	u^1	0.00421	0.00016	26

This table reports estimated parameter values for the structural model in Section 3. Panel A summarizes the parameter values that we calibrated directly using observable targets. The target moments for robo returns are calculated as the mean and standard deviation of the monthly log return within the robo-advisor, taken across all investor-month level observations in our sample. The target moments for self-directed returns are similarly calculated using monthly observations among investors who are not enrolled in the robo-advisor. Panel B summarizes the parameter values that are estimated. We use minimum distance estimation to infer these parameters based the moments illustrated in Figure 8, which are survival rates conditional on i) type of advisor, and ii) holding returns within the robo-advisor.

Table 4. Value of Complementarity and Counterfactuals

		Type 0 Advisor	Type 1 Advisor	Type 0 + more info	Type 0 + more utility
Client	Value	60.68	69.81	61.12	69.26
	Diff	0.00	9.13	0.44	8.58
	% Diff	0.00	15.05	0.72	14.14
Firm	Value	1.95	2.11	1.96	2.09
	Diff	0.00	0.16	0.01	0.14
	% Diff	0.00	8.14	0.72	7.36

This table shows values of the surplus from enrolment in robo-advice, which are defined rigorously in Appendix OA.1. The top three rows show client surplus, measured in units of expected lifetime utility, while the bottom three rows show firm surplus measured in units of expected discounted cash flows as a percentage of assets under management. The first two columns show values of surplus with a type 0 (low-retention) and type 1 (high-retention) advisor, respectively. The third column is a counterfactual scenario in which a type 0 advisor achieves the same precision of prior beliefs as a type 1 advisor, holding other parameters constant. The final column is a counterfactual scenario in which a type 0 advisor achieves the same flow utility u^1 as a type 1 advisor, holding other parameters constant.

Online Appendix

(Not for publication)

OA.1 Structural Model Derivations

OA.1.1 The Investor's Problem

Let $V^j(w, m, \tau)$ be the value function of an investor who is still enrolled in robo-advice. Let $\bar{V}(w)$ be the value of an investor who has quit and, by assumption, cannot re-enrol.

First, we guess that $\bar{V}(w) = w + \frac{\bar{r}^0}{\rho}$. This value must satisfy the Bellman equation

$$\bar{V}(w) = \rho\mathbb{E}\left[w'\right] + (1-\rho)\,\mathbb{E}\left[\bar{V}\left(w'\right)\right],$$

where next period's wealth is $w' = w + r^{0\prime}$ under self-directed investment. To verify our guess, substitute on the right-hand side to get

$$\begin{aligned}\rho\mathbb{E}\left[w'\right] + (1-\rho)\,\mathbb{E}\left[\bar{V}\left(w'\right)\right] &= \rho\left(w + \bar{r}^0\right) + (1-\rho)\left(w + \bar{r}^0 + \frac{\bar{r}^0}{\rho}\right)\\ &= w + \frac{\bar{r}^0}{\rho}.\end{aligned}$$

Second, we guess that $V^j(w, m, \tau) = w + \frac{\bar{r}^0}{\rho} + F^j(m, \tau)$. This value must satisfy the Bellman equation

$$V^j(w, m, \tau) = \max\left\{\bar{V}(w), u^j + \rho\mathbb{E}\left[w'|m, \tau\right] + (1-\rho)\,\mathbb{E}\left[V^j\left(w', m', \tau'\right)|m, \tau\right]\right\}$$

where $w' = w + r'$. We have $\mathbb{E}\left[w'|m, \tau\right] = w + \bar{r}^0 + m$. With our guess, the second term under the max operator is equal to

$$\begin{aligned}u^j + \rho\mathbb{E}\left[w'|m, \tau\right] + (1-\rho)\,\mathbb{E}\left[V^j\left(w', m', \tau'\right)|m, \tau\right] &= u^j + \rho\left(w + \bar{r}^0 + m\right) + (1-\rho)\left(w + \bar{r}^0 + m + \frac{\bar{r}^0}{\rho} + \mathbb{E}\left[F^j\left(m', \tau'\right)|m, \tau\right]\right)\\ &= w + \bar{r}^0\left(1 + \frac{1-\rho}{\rho}\right) + u^j + m + (1-\rho)\,\mathbb{E}\left[F^j\left(m', \tau'\right)|m, \tau\right]\\ &= w + \frac{\bar{r}^0}{\rho} + u^j + m + (1-\rho)\,\mathbb{E}\left[F^j\left(m', \tau'\right)|m, \tau\right]\end{aligned}$$

and the Bellman equation becomes

$$w + \frac{\bar{r}^0}{\rho} + F^j(m, \tau) = w + \frac{\bar{r}^0}{\rho} + \max\left\{0, u^j + m + (1-\rho)\,\mathbb{E}\left[F^j\left(m', \tau'\right)|m, \tau\right]\right\}$$

which verifies our guess as long as $F^j(m, \tau)$ satisfies Equation (10). For the terminal condition, notice that the distribution of m' collapses to a point mass at m as $\tau \to \infty$ and, since $\tau' = \tau + \tau_\varepsilon$, we get

$$\begin{aligned}F^j_\infty(m) \equiv \lim_{\tau\to\infty} F^j(m, \tau) &= \max\left\{0, u^j + m + (1-\rho)\lim_{\tau\to\infty} F^j(m, \tau + \tau_\varepsilon)\right\}\\ &= \max\left\{0, u^j + m + (1-\rho)\lim_{\tau\to\infty} F^j(m, \tau + \tau_\varepsilon)\right\}\\ &= \max\left\{0, u^j + m + (1-\rho)\,F^j_\infty(m)\right\}\end{aligned}$$

We guess that $F_\infty^j(m) = \frac{\max\{0, m+u^j\}}{\rho}$. If $m + u^j < 0$, the equation above is trivially satisfied since it becomes $0 = 0$. If $m + u^j > 0$, we can verify our guess by substituting on the right-hand side above to get

$$\begin{aligned}\max\left\{0, u^j + m + (1-\rho) F_\infty^j(m)\right\} &= \left(u^j + m\right)\left(1 + \frac{1-\rho}{\rho}\right) \\ &= \frac{u^j + m}{\rho}.\end{aligned}$$

as claimed in the text.

OA.1.2 Investor and Firm Surplus

Consider an investor initially enrolled with a human advisor type j. Consider a sequence of returns and preference shocks such that, if her horizon does not end before t, the investor quits exactly at date t after enrolment. In this event, her lifetime utility, conditional on the true θ, is

$$\begin{aligned}\mathbb{E}\left[U|\text{quit at } t, \theta, j\right] &= w + \frac{\bar{r}^0}{\rho} + \left(\theta + u^j\right)\left[1 + (1-\rho) + \cdots + (1-\rho)^{t-1}\right] \\ &= w + \frac{\bar{r}^0}{\rho} + \left(\theta + u^j\right)\frac{1-(1-\rho)^t}{\rho}\end{aligned}$$

By contrast the utility of an investor who only invests in her outside return is $w + \frac{\bar{r}^0}{\rho}$. Averaging across all sequences of shocks, the expected surplus from enrollment is

$$\bar{S} = \left(\theta + u^j\right)\sum_{t=1}^{\infty} Pr\left[\text{quit at } t\right] * \left(\frac{1-(1-\rho)^t}{\rho}\right)$$

For the the advisory firm, we assume that it charges a fee $F \cdot W_t$ each period until the investor either reaches the end of her investment horizon and consumes, or quits, where F denotes fees per unit of assets under management. Conditional on the true θ, the log-normal distribution of returns implies that $\mathbb{E}\left[W_t\right] = W_0 \bar{R}^t$, where $\bar{R} = \exp\left\{\theta + \frac{1}{2}\sigma_\varepsilon^2\right\}$. is the expected per-period return on wealth. The firm discounts future fee revenue at the rate γ, and we denote the present value of all fees by Π. We then have

$$\begin{aligned}\mathbb{E}\left[\Pi|\text{quit at } t, \theta, j\right] &= FW_0 + (1-\rho)\frac{\mathbb{E}\left[FW_1\right]}{1+\gamma} + \cdots + (1-\rho)^{t-1}\frac{\mathbb{E}\left[FW_{t-1}\right]}{(1+\gamma)^{t-1}} \\ &= FW_0\left[1 + \frac{1-\rho}{1+\gamma}\bar{R} + \cdots + \left(\frac{1-\rho}{1+\gamma}\bar{R}\right)^{t-1}\right] \\ &= FW_0\frac{1-(1-\chi)^t}{\chi}\end{aligned}$$

where we have defined $\chi = 1 - \frac{1-\rho}{1+\gamma}\bar{R}$. Averaging across all sequences of shocks, the expected firm

surplus from enrollment is therefore

$$\bar{\Pi} = FW_0 \sum_{t=1}^{\infty} Pr\left[\text{quit at } t\right] * \left(\frac{1-(1-\chi)^t}{\chi}\right)$$

In our empirical results in Table 4, we report $\bar{\Pi}/W_0$, i.e., the firm's surplus per unit of initial assets under management.

OA.2 Numerical Solution Method

OA.2.1 Solving for Value Functions

For our numerical solution, we change the state variable in the investor's problem to $\tilde{m} = m + w^j$. The change of variable from m to $\tilde{m}$ is motivated by the fact that the numerical solution of this equation does not depend on any of the estimated parameters of the model. Hence, we only solve the Bellman equation once, and then look up the values implied by different parameter constellations during the estimation procedure. Repeating the argument in Appendix OA.2.1, we obtain the following Bellman equations:

$$\begin{aligned} V\left(w, \tilde{m}, \tau\right) &= w + \frac{\theta_0}{\rho} + F\left(\tilde{m}, \tau\right) \\ F\left(\tilde{m}, \tau\right) &= \max\left\{0, \tilde{m} + (1-\rho)\,\mathbb{E}\left[F\left(\tilde{m}', \tau'\right)|\tilde{m}, \tau\right]\right\} \end{aligned}$$

The belief updating equations, with this change of variable and in recursive notation, become

$$\begin{aligned} \tilde{m}' &= \tilde{m} + \frac{\tau_\varepsilon}{\tau + \tau_\varepsilon}\left(r^{e,\prime} + w^j - \tilde{m}\right) \\ \tau' &= \tau + \tau_\varepsilon \end{aligned}$$

Note that $\mathbb{E}\left[r' + w^j|\tilde{m}, \tau\right] = \tilde{m}$ and

$$\begin{aligned} \mathbb{V}\text{ar}\left[r'|\tilde{m}, \tau\right] &= \mathbb{V}\text{ar}\left[\theta + \varepsilon'|\tilde{m}, \tau\right] \\ &= \frac{1}{\tau} + \frac{1}{\tau_\varepsilon} \end{aligned}$$

Hence, the conditional probability distribution of next period's state is given by

$$\tilde{m}'|\tilde{m}, \tau \sim N\left(\tilde{m}, v\left(\tau\right)\right),$$

where

$$v\left(\tau\right) = \left(\frac{\tau_\varepsilon}{\tau + \tau_\varepsilon}\right)^2 \left(\frac{1}{\tau} + \frac{1}{\tau_\varepsilon}\right)$$

In order to solve the Bellman equation, we therefore solve

$$F\left(\tilde{m}, \tau\right) = \max\left\{0, \tilde{m} + (1-\rho)\int F\left(\tilde{m}', \tau + \tau_\varepsilon\right) d\Phi\left(\frac{\tilde{m}' - \tilde{m}}{\sqrt{v\left(\tau\right)}}\right)\right\}$$

with terminal condition $\lim_{\tau\to\infty} F(\tilde{m},\tau) = \frac{\max\{0,\tilde{m}\}}{\rho}$, where $\Phi(.)$ is the standard normal CDF. We describe our grid-based approximation of this equation below.

Given this solution, we also obtain the continuation value $C(\tilde{m},\tau)$, which are equal to the second term under the max operator in the above expression. We can then also compute the state-contingent hazard rates

$$\Pr[\text{quit}|m,\tau,\text{enrolled}] \equiv \lambda(\tilde{m},\tau) = \frac{1}{1+\exp\left\{-\frac{1}{\eta}C(\tilde{m},\tau)\right\}}$$

which follow from our assumption in Equation (12).

OA.2.2 Solving for Quit Probabilities

In order to calculate model-implied quit probabilities without simulation, we define the probability

$$U_t^j(\tilde{m}) = \Pr[\text{enrolled at } t, \tilde{m}|\theta, j]$$

Conditional on the true θ and assignment to human advisor type j, this is the joint probability density for the event that a client is still enrolled at time t (i.e., does not quit at any previous time $s = 1,\ldots,t-1$) and faces current state variable $\tilde{m}$. Notice that, for a given parameter constellation, the precision of beliefs at time t is deterministic and equal to $\tau_t^j = \tau_0^j + t\cdot\tau_\varepsilon$. Since the investor does not quit at date 0 (by assumption), we have the initialization

$$U_1^j(\tilde{m}) = \Pr[\tilde{m}_1 = \tilde{m}|\theta, j]$$

The joint probabilities at subsequent dates satisfy the Kolmogorov forward equation

$$U_{t+1}^j(\tilde{m}') = \int U_t^j(\tilde{m})\left(1-\lambda\left(\tilde{m},\tau_t^j\right)\right)P\left(d\tilde{m},\tilde{m}'\right)$$

where $h(\tilde{m},\tau)$ is the conditional hazard rate defined in the previous section, and where $P(\tilde{m},\tilde{m}')$ is a kernel describing the transition probabilities from state $\tilde{m}_t = \tilde{m}$ to state $\tilde{m}_{t+1} = \tilde{m}'$. In order to compute this kernel, we again use the updating equation

$$\tilde{m}' = \tilde{m} + \frac{\tau_\varepsilon}{\tau_t^j+\tau_\varepsilon}\left(r_{t+1}^e + u^j - \tilde{m}\right)$$

Note that, conditional on the true θ, $\mathbb{E}[r_{t+1}|\theta] = \theta$ and $\mathbb{V}\text{ar}[r_{t+1}|\theta] = \frac{1}{\tau_\varepsilon}$. Hence,

$$\tilde{m}'|\tilde{m},\theta,t \sim N\left(\mu_\theta(\tilde{m}), v_\theta\left(\tau_t^j\right)\right)$$

where

$$\mu_\theta(\tilde{m}) = \tilde{m} + \frac{\tau_\varepsilon}{\tau_t^j+\tau_\varepsilon}\left(\theta + u^j - \tilde{m}\right)$$

$$v_\theta\left(\tau_t^j\right) = \left(\frac{\tau_\varepsilon}{\tau_t^j+\tau_\varepsilon}\right)^2\frac{1}{\tau_\varepsilon}$$

Combining, the Kolmogorov forward equation becomes

$$U_{t+1}^{j}\left(\tilde{m}'\right)=\int U_{t}^{j}\left(\tilde{m}\right)\left(1-h\left(\tilde{m},\tau\right)\right)d\Phi\left(\frac{\tilde{m}'-\mu_{\theta}\left(\tilde{m}\right)}{\sqrt{\upsilon_{\theta}\left(\tau_{t}^{j}\right)}}\right)$$

where $\Phi\left(.\right)$ is the standard normal CDF. In order to obtain the model-implied Kaplan-Meier survival curves, used for structural estimation, we define the average holding return up to time t as $\bar{r}_{t}=\frac{1}{t}\sum_{s=1}^{t}r_{t}$, and define the average excess return $\bar{r}_{t}^{e}=\bar{r}_{t}-\bar{r}^{0}$. We note that the Kalman filter in Equation (7) is equivalent to

$$\tilde{m}_{t}^{j}=w^{j}+m_{0}^{j}+\frac{t\tau_{\varepsilon}}{t\tau_{\varepsilon}+\tau_{0}^{j}}\left[\bar{r}_{t}^{e}-m_{0}^{j}\right]\equiv\mu_{t}^{j}\left(\bar{r}_{t}\right)$$

Hence a bucket $\bar{r}_{t}\in b\equiv\left[b_{0},b_{1}\right]$, is equivalent to $\tilde{m}_{t}^{j}\in\left[\mu_{t}^{j}\left(b_{0}\right),\mu_{t}^{j}\left(b_{1}\right)\right]$. We then compute the bucket-conditional hazard rate as

$$\begin{aligned}Pr\left[\text{quit at }t|\text{enrolled at }t,\bar{r}_{t}\in\left[b_{0},b_{1}\right],\theta,j\right]&=\frac{\int_{\mu_{t}^{j}\left(b_{0}\right)}^{\mu_{t}^{j}\left(b_{1}\right)}U_{t}^{j}\left(\tilde{m}\right)h\left(\tilde{m},\tau_{t}^{j}\right)d\tilde{m}}{\int_{\mu_{t}^{j}\left(b_{0}\right)}^{\mu_{t}^{j}\left(b_{1}\right)}U_{t}^{j}\left(\tilde{m}\right)d\tilde{m}}\\&\equiv\lambda_{t}^{j}\left(b\right)\end{aligned}$$

Finally we compute

$$g_{j,b,t}(\alpha)=\prod_{s\leq t}\left(1-\lambda_{t}^{j}\left(b\right)\right)$$

which is the model-implied Kaplan-Meier survival curve conditional on the return bucket b.

OA.2.3 Discretized State Space

We compute the Bellman and Kolmogorov equations outlined above, as well as bucket-conditional hazard rates, using a discretized space for the investor's state variables $(\tilde{m},\tau)$. Given calibrated parameters, we consider a 1000-point grid for $\tau\in\tau_{\varepsilon}*\{0,1,\ldots,1000\}$, where $\tau_{\varepsilon}=1/\sigma_{\varepsilon}^{2}$ is the precision of return shocks, and a 500-point grid for $\tilde{m}\in[\theta-3*\sigma_{\varepsilon},\theta+3*\sigma_{\varepsilon}]$ to compute the Bellman equation. For computational speed, when solving the Kolmogorov equation for a particular parameter constellation, we specialize to a smaller 100-point grid for $\tilde{m}\in\left[\tilde{m}_{0}^{j}-\frac{2}{\sqrt{\tau_{0}^{j}}},\tilde{m}_{0}^{j}+\frac{2}{\sqrt{\tau_{0}^{j}}}\right]$, where $\tilde{m}_{0}^{j}=\theta+w^{j}$ is the initial value of $\tilde{m}$.

We obtain the transition probabilities from $\tilde{m}$ to $\tilde{m}'$ on this grid using the Tauchen (1986) method for grid approximation. In general, for any continuous random variable x, approximated by a grid $\hat{x}=\left\{x^{(1)},\ldots,x^{(N)}\right\}$, we use the following discretized probabilities:

$$P\left(x^{(k)}\right)=\begin{cases}Pr\left[x\leq x^{(1)}+\frac{d}{2}\right], & k=1\\Pr\left[x\in\left(x^{(k)}-\frac{d}{2},x^{(k)}+\frac{d}{2}\right]\right], & k=1,\ldots N_{m}-1\\Pr\left[x>m^{(N)}-\frac{d}{2}\right], & k=N_{m}\end{cases}$$

where d is the distance between two adjacent points on the grid. If, as in our application, x has a Gaussian distribution with mean μ and standard deviation s, we have

$$P\left(x^{(k)};\mu,s\right)=\begin{cases}\Phi\left(\frac{x^{(1)}+\frac{d}{2}-\mu}{s}\right), & k=1\\ \Phi\left(\frac{x^{(k)}+\frac{d}{2}-\mu}{s}\right)-\Phi\left(\frac{x^{(k)}-\frac{d}{2}-\mu}{s}\right), & k=2,\dots N-1\\ 1-\Phi\left(\frac{x^{(N)}-\frac{d}{2}-\mu}{s}\right), & k=N\end{cases} \tag{13}$$

where $\Phi(.)$ is the standard Gaussian cumulative distribution function.

OA.3 Structural Estimation and Inference

OA.3.1 Minimum Distance Estimation

The canonical minimum distance estimation problem (see Cameron and Trivedi (2005), Chapter 6.7) is

$$\hat{\alpha}=\arg\max_{\alpha}\left(\hat{\pi}-g(\alpha)\right)'W\left(\hat{\pi}-g(\alpha)\right)$$

where α is a vector of q structural parameters, $\hat{\pi}$ is a vector of $r>q$ reduced form parameters from the data, and $g(\alpha)$ are the model-implied reduced form parameters/moments. In our case we have structural parameters $\alpha=\left(\eta,\left(\tau^j,u^j\right)_{j\in\{0,1\}}\right)$ so that $q=5$ and $\hat{\pi}=(\hat{\pi}_{j,b,t})$ denotes the vector of Kaplan-Meier survival estimates in the data for advisor type $j\in\{0,1\}$, return buckets $b\in\{\text{High},\text{Low}\}$ and time periods $t=1,\dots,15$ since enrollment, so that $r=60$. We set the weighting matrix $W=I$ to the identity matrix, which yields the objective function in Equation (11).

For our numerical procedure, we further exploit the fact that model-implied moments for advisor type j depend only on i) the noise parameter, and ii) the precision and flow utility parameters $\left(\tau^j,u^j\right)$ for type j. Thus, we can re-write the estimation problem as a two-step problem:

$$\hat{\alpha}=\arg\max_{\eta}\sum_j\max_{(\tau^j,u^j)}\sum_b\sum_t(\hat{\pi}_{j,b,t}-g_{j,b,t}(\eta,\tau^j,u^j))^2 \tag{14}$$

In this problem, for every noise parameter η and for every advisor type j, we can separately solve the inner maximization problem of finding parameters $\left(\tau^j,u^j\right)$ that minimize the type-specific distance between empirical and model-implied moments. The outer maximization problem is to find the noise parameter that implies the best overall fit across advisor types.

Our estimation routine can be summarized as follows:

- Set calibrated parameters, as reported in Table 3 Panel A. Set a grid for the noise parameter $\eta\in\{\eta_1,\dots\eta_M\}$.
- Solve the Bellman Equation using the change of variable $\tilde{m}=m+u^j$, as outlined in Appendix OA.2.1, so that the solution is independent of α
- For every grid point $m=1,\dots,M$ and every advisor type $j\in\{0,1\}$:
  - Fix the noise parameter $\eta=\eta_m$

– Solve the inner maximization problem over $\left(\tau^j, u^j\right)$ in Equation (14) using the Nelder-Mead simplex algorithm (using Matlab's *fminsearch* function). At every step of this algorithm, the objective function is evaluated by solving the Kolmogorov forward equation, as outlined in Appendix OA.2.2, to get model-implied quit probabilities $g_{j,b,t}(\eta_m, \tau^j, u^j)$

– Store the maximizing values $\left(\tau_m^j, u_m^j\right)$ and the maximized objective

$$v^j(m) = \sum_b \sum_t (\hat{\pi}_{j,b,t} - g_{j,b,t}(\eta_m, \tau_m^j, u^j))^2$$

- Choose the best grid fitting point $m^\star = \arg\min \sum_j v^j(m)$. Set the estimated noise parameter to $\eta = \eta_{m^\star}$ and the remaining estimated parameters to $\left(\tau_{m^\star}^j, u_{m^\star}^j\right)$, which are reported in Table 3 Panel B.

Figure OA.I below plots the maximizing values $\left(\tau_m^j, u_m^j\right)$ from the inner loop as a function of the noise parameter on grid points η_m. The plots show that the qualitative properties of our estimated parameters, and the ranking across advisor types, do not depend on a specific choice of η.

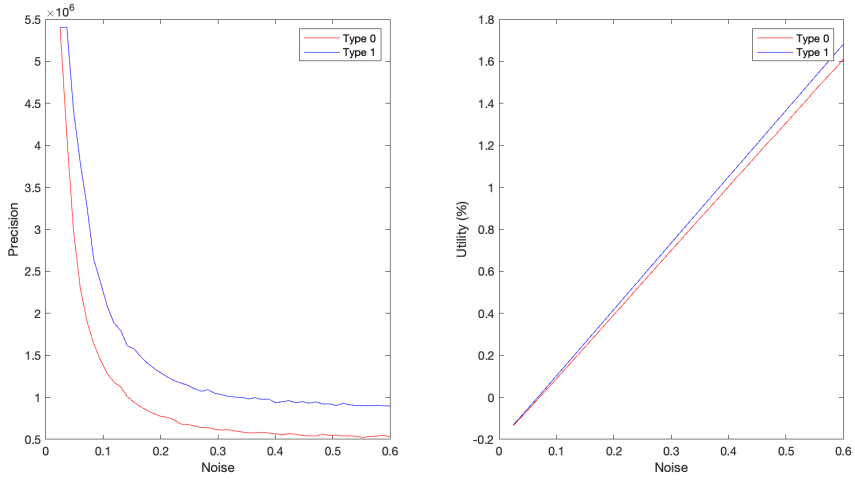


Figure OA.I: This figure plots the maximizing values $\left(\tau_m^j, u_m^j\right)$ from the inner loop as a function of the noise parameter on grid points η_m

OA.3.2 Bootstrap Procedure

Our bootstrap procedure involves repeatedly sampling observations from the original dataset with replacement to create multiple bootstrap samples indexed by $n = 1, \ldots, B$, where we set $B = 1000$. Bootstrap samples are used to estimate the variance-covariance matrix of the 60 reduced form parameters $\hat{\pi}$, which are the Kaplan-Meier survival estimates illustrated in Figure 8.

We conducted two different types of bootstrap draws. In the first, we simply draw observations (investor-date indicators of survival) at random with replacement from the underlying data, with sample size equal to the size of the initial dataset. In the second, following Politis and Romano (1994)

we draw time-periods with replacement, and for each draw of a time-period, we employ all associated investor observations. This procedure controls for the possibility of common shocks that affect all investors at a given point in time. As the number of investors may differ across time periods, this second procedure can produce datasets of a total size different to the original dataset, but with the same time-series length. In the paper, we show the results from using the second, clustered bootstrap, as the standard errors that we estimate using this procedure are larger and therefore more conservative.

For every bootstrap sample n, we generate separate Kaplan-Meier estimates $\hat{\pi}_{j,b,t}^{(n)}$ for advisor types $j \in \{0,1\}$, return buckets $b \in \{\text{High}, \text{Low}\}$ and $t = 1, \ldots, 15$ and collect them in the 60×1 vector $\hat{\pi}^{(n)}$. We then calculate the sample variance-covariance matrix $\hat{\Omega}$ of the reduced form parameters, which has dimension 60×60, as follows:

$$\begin{aligned}\hat{\Omega} &= \hat{\mathbb{V}\text{ar}}\left[\hat{\pi}\right] \\ &= \frac{1}{B-1}\sum_{n=1}^{B}\left(\hat{\pi}^{(n)} - \bar{\pi}\right)\left(\hat{\pi}^{(n)} - \bar{\pi}\right)'\end{aligned}$$

where $\bar{\pi} = \frac{1}{B}\sum_{n=1}^{B}\hat{\pi}^{(n)}$ is the mean across bootstrapped samples.

OA.3.3 Standard Errors

The standard formula for the asymptotic variance-covariance matrix of minimum distance estimators $\hat{\alpha}$ is given by

$$\hat{\mathbb{V}\text{ar}}\left[\hat{\alpha}\right] = (\hat{G}'W\hat{G})^{-1}\hat{G}'W\hat{\Omega}W\hat{G}(\hat{G}'W\hat{G})^{-1}$$

where $\hat{\Omega}$ is the variance-covariance matrix of the reduced-form parameters/moments $\hat{\pi}$, which we obtain using the bootstrap procedure outlined above, W is the weighting matrix in the objective function, which we set to $W = I$, and $\hat{G} = \frac{\partial g(\hat{\alpha})}{\partial \alpha}$ is the $r \times q$ Jacobian matrix of local derivatives of the $r = 60$ model-implied survival probabilities to the $q = 5$ parameters.

In our setting, let $g_j\left(\eta, \tau^j, u^j\right)$ denote the 30×1 vector of stacked model-implied survival probabilities for advisor type j (i.e., probabilities $g_{j,b,t}\left(\eta, \tau^j, u^j\right)$ for two return buckets b and 15 times periods t). The Jacobian matrix we compute has the following block structure:

$$\hat{G} = \begin{pmatrix} \frac{\partial g_1(\eta,\tau^1,u^1)}{\partial(\tau^1,u^1)} & 0 & \frac{\partial g_1(\eta,\tau^1,u^1)}{\partial \eta} \\ 0 & \frac{\partial g_2(\eta,\tau^2,u^2)}{\partial(\tau^2,u^2)} & \frac{\partial g_2(\eta,\tau^2,u^2)}{\partial \eta} \end{pmatrix}$$

where $\frac{\partial g_1(\eta,\tau^1,u^1)}{\partial(\tau^1,u^1)}$ and $\frac{\partial g_2(\eta,\tau^2,u^2)}{\partial(\tau^2,u^2)}$ are sub-matrices of gradients with respect to type-specific parameters, each with dimension 30×2, and $\frac{\partial g(\eta,\tau^1,u^1)}{\partial \eta}$ and $\frac{\partial g(\eta,\tau^1,u^1)}{\partial \eta}$ are sub-matrices of gradients with respect to the fitting/noise parameter, each with dimension 30×1.

We compute gradients numerically by using a two-sided approximation, for example, for the noise parameter:

$$\frac{\partial g_1(\eta,\tau^1,u^1)}{\partial \eta} = \frac{g_1(\eta+h,\tau^1,u^1) - g_1(\eta-h,\tau^1,u^1)}{2h}$$

We then evaluate $\hat{\mathbb{V}\text{ar}}[\hat{\alpha}]$ and report the 5 estimated standard errors of parameters $\hat{\alpha}$ as follows:

$$\hat{SE}[\hat{\alpha}] = \sqrt{\text{diag}\left(\hat{\mathbb{V}\text{ar}}[\hat{\alpha}]\right)}$$

and calculate the associated t-statistics as $t = \hat{\alpha}/\hat{SE}[\hat{\alpha}]$.

OA.4 Discussion of Structural Identification

OA.4.1 Approximate Closed-Form Solution

We first outline an approximate version of investors' optimization problem in order to provide analytical insights into the sources of structural identification. Assume that each period, the investor decides whether or not to quit, as in the full model, but treats this decision as if it was final. In other words, if she does not quit, she intends to remain enrolled until the end of her investment horizon. In this case, it is easy to see that the perceived continuation value from staying enrolled is

$$\hat{C}^j(m) = \frac{m + u^j}{\rho}$$

Intuitively, the investor enjoys an expected flow of future returns and utility of m per period, and stays enrolled until the end of her investment horizon, which is on average $\frac{1}{\rho}$ periods. Using Equation (7), beliefs after t periods of enrollment can be written as

$$m_t^j = m_0^j + \frac{t\tau_\varepsilon}{t\tau_\varepsilon + \tau_0^j}\left(\bar{r}_t^e - m_0^j\right)$$

where $\bar{r}_t^e = \sum_{s \le t} r_s - \bar{r}^0$ is the excess log holding return up to date t.

With the logistic choice rule described in Equation (12), the log odds ratio of quitting at date t is given by

$$\begin{aligned}\lambda_t(\bar{r}_t) \equiv &= -\frac{1}{\eta}\hat{C}^j(m_t) \\ &= \frac{1}{\eta\rho}\left[m_0^j + \frac{t\tau_\varepsilon}{t\tau_\varepsilon + \tau_0^j}\left(\bar{r}_t^e - m_0^j\right) + u^j\right]\end{aligned}$$

In this equation, ρ and τ_ε and m_0^j are parameters that we calibrate in Table 3 Panel A, and the parameters that we estimate are the noise parameter η, the initial precision of beliefs τ_0^j, and the utility parameter u^j.

Clearly, the odds of quitting are increasing in the utility parameter u^j, so that the level of quit rates helps us to identify flow utilities. Another moment we implicitly use for identification is the sensitivity of quit probabilities to returns, which relates to

$$\frac{\partial \lambda_t}{\partial \bar{r}_t} = \frac{1}{\eta\rho}\frac{t\tau_\varepsilon}{\tau_0^j + t\tau_\varepsilon}$$

and therefore depends on both σ and τ_0^j. However it is clear that this moment places greater weight

on τ_0^j in early periods, and close to zero weight on τ_0^j in late periods (e.g., $\frac{t\tau_\varepsilon}{\tau_0^j + t\tau_\varepsilon} \to 1$ regardless of τ_0^j when $t \to \infty$) . Thus, the mapping between sensitivities at different time periods and parameters $\left(\eta, \tau_0^j\right)$ can be inverted. Intuitively, sensitivity in later periods identify η in isolation, while the early moments identify a combination of τ_0^j and η. Although we do not attempt to estimate the prior means μ_0^j separately, a similar argument could be used in principle to identify this parameter separately from w^j.

OA.4.2 Empirical Sensitivities of Targets to Parameters

The following figures show how the model-implied survival probabilities at different horizons vary as a function of the four parameters, namely, σ_ε, τ_0^0, u^0, τ_0^1, and u^1. Together, the figures provide visual evidence of the differential sensitivity of moments to model parameters.

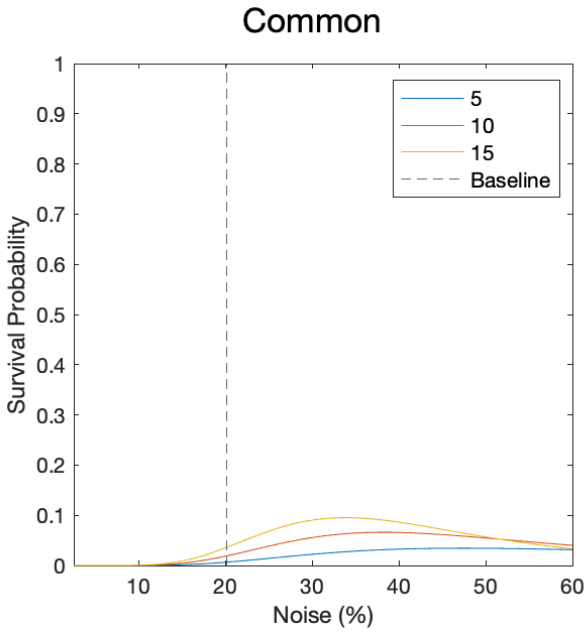


Figure OA.II: This figure illustrates how the model-implied survival probability at different periods (5, 10, and 15 months) varies as we vary σ_ε, the volatility of returns conditional on θ, i.e., the extent of "noise in the signal". This parameter is common to all investors in the model regardless of the advisor type to which they are assigned.

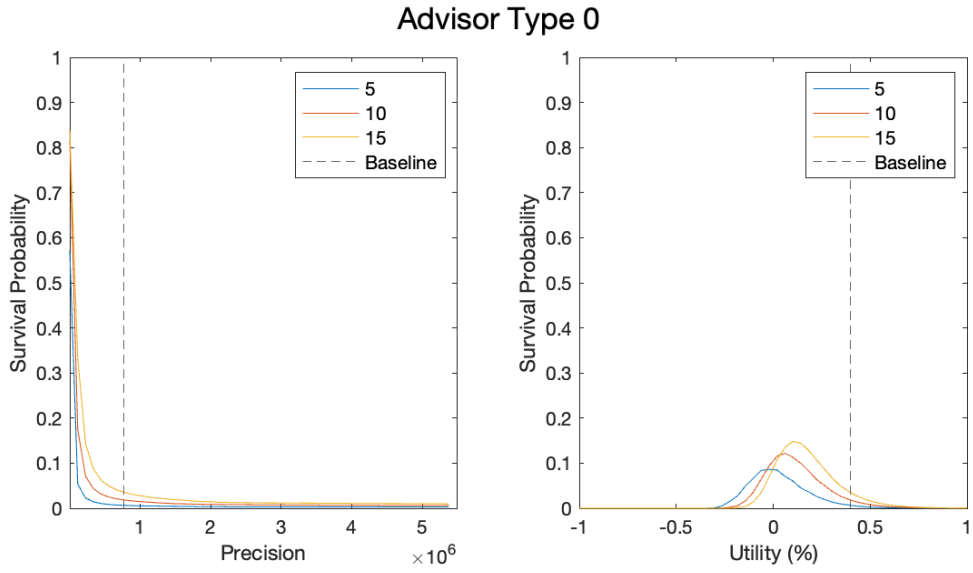


Figure OA.III: This figure illustrates how the model-implied survival probability at different periods (5, 10, and 15 months) for investors assigned to advisors of type 0 varies as we vary the prior precision τ_0^0 (left panel), and the per-period utility u^0 (right panel).

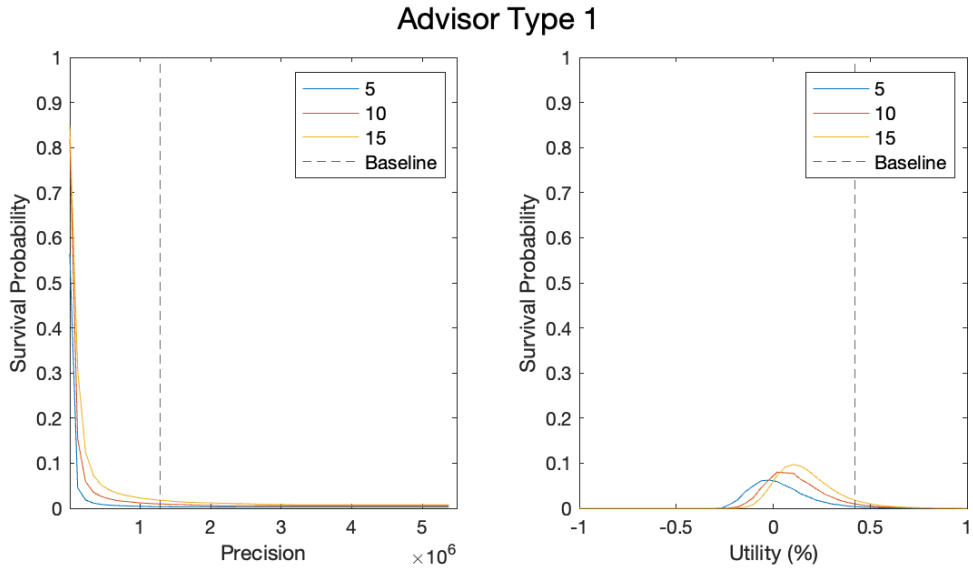


Figure OA.IV: This figure illustrates how the model-implied survival probability at different periods (5, 10, and 15 months) for investors assigned to advisors of type 1 varies as we vary the prior precision τ_0^1 (left panel), and the per-period utility u^1 (right panel).

OA.5 Additional Figures and Tables

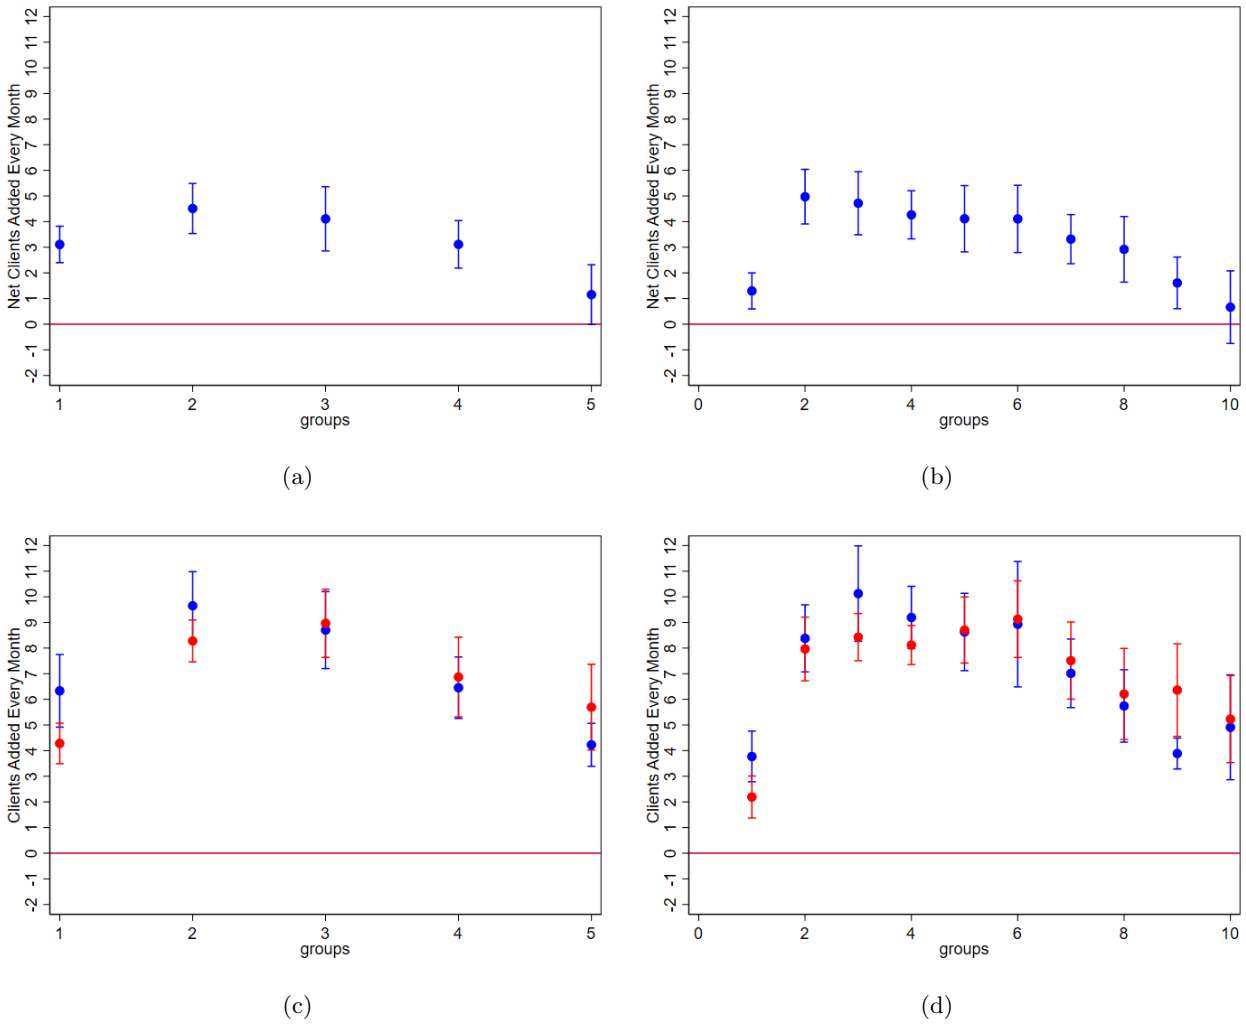


Figure OA.V: Subfigure (a) is constructed by computing, for every advisor, the number of investors they advise at the beginning of each month and sorting all the available advisors into quintiles based on their current workload. The figure then reports the average net increase in the number of clients allocated to advisors in each group every month, computed as the number of investors allocated to each advisor minus the number of investors lost by each advisor every month because of attrition, together with 95% confidence intervals. Subfigure (b) repeats the analysis using deciles instead of quintiles. Subfigures (c) and (d) repeat the exercise but focus only on investors' additions and split the advisors into high-retention (blue) and low-retention (red).

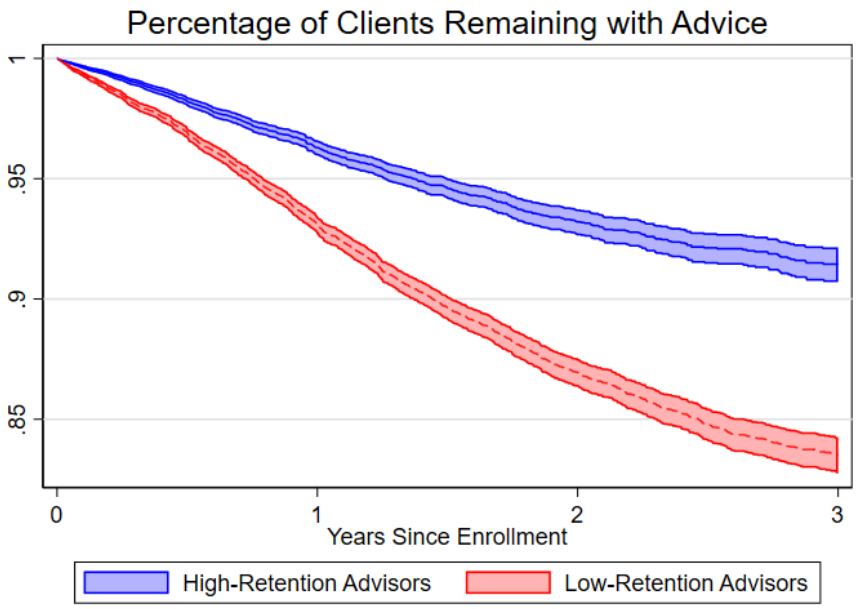


Figure OA.VI: This figure plots Kaplan-Meier survival functions for investors assigned to high- and low-retention advisors, excluding clients who interacted with multiple advisors during their tenure with the robo-advisor. Survival rates for clients assigned to high-retention advisors are shown in blue; survival rates for clients assigned to low-retention advisors are shown in red.

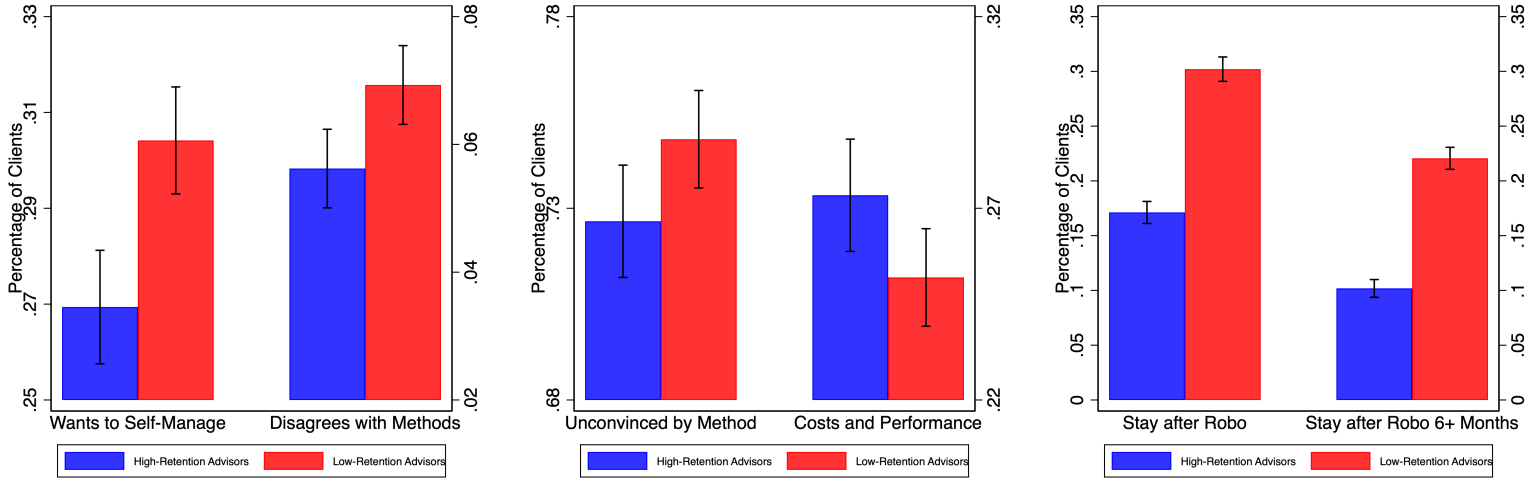


Figure OA.VII: The left subfigure reports the percentage of clients quitting robo-advice because they want to self-manage their portfolio (first set of bars) and because they disagree with the robo-advisors' investment methodology (second set of bars), together with 95% confidence intervals. The middle subfigure collects the 43 reasons for quitting robo-advice, and classifies them into reasons related to investors wanting to self-manage their portfolio (first set of bars) and reasons related to robo-advisor's costs and performance (right set of bars)— excluding all other reasons from the analysis. The right subfigure computes the percentage of clients who continue to stay with the asset manager after quitting robo-advice as self-directed investors (left set of bars) and the percentage of clients who continue to stay with the asset manager after quitting robo-advice as self-directed investors for at least 6 months (right set of bars). Blue bars denote investors assigned to high-retention advisors, and the red bars denote investors assigned to low-retention advisors.

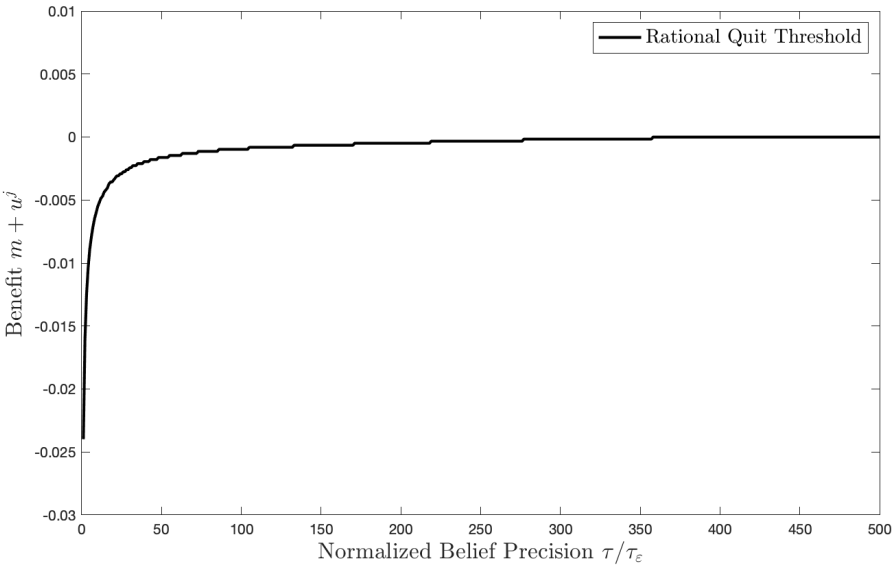


Figure OA.VIII: This figure illustrates the optimal behavior of rational investors in our structural model. As a function of the precision of the investor's beliefs τ, normalized by the precision of returns τ_ε, the solid line shows the minimum value of $m + u^j$, the sum of the investor's expected log excess return and flow utility, for which she is willing to remain enrolled in robo-advice. For our empirical application, we assume that quit decisions are based on this threshold but subject to cross-sectional noise, as defined by the logistic choice rule in Equation (12).

Table OA.I. Demographic and Portfolio Characteristics of Advised Investors

Panel A. Demographic Characteristics

	N	mean	sd	p25	p50	p75
Age	54,325	63.816	12.071	57.000	65.000	72.000
Male	54,744	0.598	0.490	0.000	1.000	1.000
Tenure	54,744	13.481	9.084	3.833	13.667	20.250

Panel B. Portfolio-Related Characteristics

	N	mean	sd	p25	p50	p75
Wealth	54,744	$758,378	$821,029	$210,800	$478,929	$981,330
NumAssets	54,744	7.952	4.910	5.000	6.000	9.000
PctAMProducts	54,744	0.974	0.069	1.000	1.000	1.000

Panel C. Asset Allocation Characteristics

	N	mean	sd	p25	p50	p75
PctMutualFunds	54,744	0.952	0.102	0.960	1.000	1.000
PctCash	54,744	0.018	0.046	0.000	0.000	0.008
PctETF	54,744	0.008	0.030	0.000	0.000	0.000
PctStocks	54,744	0.014	0.045	0.000	0.000	0.000
PctBonds	54,744	0.000	0.002	0.000	0.000	0.000

Panel D. Characteristics of Mutual Funds Held

	N	mean	sd	p25	p50	p75
AcctIndex	54,744	0.828	0.178	0.745	0.858	1.000
MgtFee	54,717	0.072	0.024	0.059	0.064	0.075
ExpRatio	54,707	0.093	0.027	0.078	0.083	0.096
TurnRatio	54,685	0.268	0.120	0.190	0.280	0.337

This table reports the demographic characteristics and portfolio allocation behavior of investors 12 months after signing up for advice. The results are computed at the investor level and include all account types, that is, taxable and non-taxable (IRA) accounts. Panel A reports demographic characteristics, Panel B focuses on portfolio-related characteristics, Panel C focuses on asset allocation characteristics, and Panel D focuses on the characteristics of the mutual funds held. For each variable, we report the number of accounts used in the computations, the mean, the standard deviation, and the 25^{th}, 50^{th}, and 75^{th} percentiles of the distribution.