Who Benefits from Robo-advising? Evidence from Machine Learning

Alberto G. Rossi*
Georgetown University

Stephen Utkus
Vanguard

September 26, 2020

Abstract

We study the effects of a large U.S. hybrid robo-adviser on the portfolios of previously self-directed investors. Across all investors, robo-advising reduces idiosyncratic risk by lowering the holdings of individual stocks and active mutual funds and raising exposure to low-cost indexed mutual funds. It further eliminates investors' home bias and increases investors' overall risk-adjusted performance, mainly by lowering investors' portfolio risk. We use a machine learning algorithm, known as Boosted Regression Trees (BRT), to explain the cross-sectional variation in the effects of advice on portfolio allocations and performance. Finally, we study the determinants of investors' sign-up and attrition.

JEL classification: D14, G11, O33

Keywords: FinTech, Portfolio Choice, Machine Learning, Individual Investors, Financial Literacy, Technology Adoption.

The paper has benefited from the comments made at presentations at the NBER Summer Institute on Aging, EFA Annual Meeting, ABFER 7th Annual Conference, Georgetown Global Virtual Seminar on FinTech, Workshop on FinTech and Machine Learning (Singapore), Vanguard Research Symposium, Cornell University, Georgetown University Vanderbilt University, University of Notre Dame, Boston University, Lehigh University, Baruch College, Michigan State University, Temple University, Vienna Graduate School of Finance, American University. We are grateful to Reena Aggarwal, Linda Allen, John Ameriks, Turan Bali, Gurdip Bakshi, Brad Barber, Agostino Capponi, James Choi, Youngmin Choi, Francesco D'Acunto, Xi Dong, Paul Gao, Xiaohui Gao, Christopher Hansman (the EFA discussant), Kathleen Hanley, Steven Kou, Craig Lewis, Claire Liang, Vladimir Mukharlyamov, Nandu Nayar, Lalitha Naveen, Andreas Neuhierl, Maureen O'Hara, Cindy Pagliaro, Lin Peng, Alessandro Previtero, Jonathan Reuter, Samuel Rosen, Jesus Salas, Gustavo Schwenkler, Berk Sensoy, Yao Shen, Vicky Tang, Sterling Yan, Bart Zhou Yueshen (the ABFER discussant), YinYin Yu, Kate Waldock, Yajun Wang, Steve Zeldes, Haibei Zhao, Dexin Zhou, Jonathan Zinman, for comments and suggestions.

*McDonough School of Business, Georgetown University, Washington, DC, USA. Fellow of the Luohan Academy. e-Mail: agr60@georgetown.edu.

Robo-advisers have surged in popularity in recent years as investors seek low-cost, automated investment advice. They allow investors to set up customized, diversified portfolios and can give access to other wealth management services previously limited to affluent investors, such as portfolio tax efficiency, cash flow forecasting, and retirement income planning. Many such services emphasize investment in low-cost index funds, minimal trading, tax efficiency and global diversification—arguably the benchmark for optimal portfolio advice from the empirical finance literature. In addition to being comparably inexpensive, robo-advisers have the potential to be superior to human financial advisers, as the latter have been shown to display behavioral biases and cognitive limitations (see Linnainmaa, Melzer, and Previtero, forthcoming). As a result, robo-advisers are quickly attracting attention from policymakers and investors at all levels.

This paper provides a comprehensive analysis of a major U.S. hybrid robo-adviser.[1] We study the evolution of investment advice among a sample of more than 80,000 individuals who were previously self-directed investors and sign up for professional financial advice during the 2015-2018 period. Our sample consists of investors with considerable portfolio wealth (median portfolio wealth is $282,000) and a willingness to take equity risk (median equity share 59%). Our main focus is on the portfolio effects of advice, including portfolio composition and risk-adjusted returns. Investors may derive other sources of value from advisory services. For example, substantial value to investors may arise from financial planning elements unrelated to portfolio construction, such as cash flow or retirement income planning or benefits plan optimization. Investors may also derive other benefits from advice, such as time delegation (see Kim, Maurer, and Mitchell, 2016), financial education/literacy benefits, or emotional/hedonic rewards such as improved financial well-being and peace of mind (see Rossi and Utkus, 2019).

The advice service in our study provides personalized investment portfolios for investors at low cost, relying principally on low-fee index mutual funds and a separate advisory fee of 0.30% or less.[2] At sign-up, investors are profiled on the basis of their financial objectives, risk-tolerance, investment horizons and demographic characteristics. They are then proposed a comprehensive financial plan,

[1]Although robo-advisers have been traditionally characterized as all-digital automated services, in recent years many have emerged with a hybrid model, combining both a highly automated investment and planning process with human contact. This paper studies the Personal Advisor Service (PAS) from Vanguard. As of February 2019, according to the website roboadvisorpros.com, PAS is the largest hybrid robo-adviser in the world with $115 billion in assets under management (AUM). Other robo-advisors, with a hybrid element or not, include Schwab Intelligent Portfolios ($37 billion), Betterment ($16 billion), Wealthfront ($11 billion), and Personal Capital ($8.5 billion).

[2]Fees are 0.30% on assets below $5 million; 0.2% on assets from $5 million to below $10 million; 0.1% on assets from $10 million to below $25 million; and 0.05% on assets of $25 million and above

which includes a cash flow forecast, a probability of success in achieving goals (such as financing a secure retirement), and a recommended portfolio strategy designed to help achieve such goals. The advisory algorithm maps investors to one of five risk glide-paths depending on risk tolerance and goal time horizon. Before signing-up, investors interact with a human adviser who describes the plan and may adjust it (subject to standardized guardrails) based on investor feedback. Investors are officially enrolled into the service only after accepting the proposed (or revised) plan and agreeing to move forward with the engagement. From that moment, trading occurs automatically on behalf of the investor to reach the desired portfolio allocation. Investor positions are revisited quarterly by the algorithm and trades are placed if portfolio weights deviate substantially from target weights.

We first explore the effects of robo-advice across our entire sample of investors which, as noted previously, consists of individuals who previously made their investment decisions on a self-directed basis. Advice operates significant changes on investors' portfolios. The percentage of wealth in indexed mutual funds almost doubles: it increases from 47% to 83%. Investors' international diversification increases threefold: the percentage of wealth in international mutual funds increases from 10% to 32%. Moving investors into indexed mutual funds translates into lower fees: average expense ratios are more than halved—from 19 to 9 basis points.

Advice also changes investors' risky share. It increases investors' bond-holdings from 24% to 40% and decreases investors' cash and money market mutual fund holdings from 22% to 1%. Equity holdings increase, on average, from 54% to 59% and we find large positive changes (in excess of 30%) for investors that had equity shares below 10% and negative changes (close to 30%) for those who had equity shares above 90%.

Advice also operates very large changes on the investment vehicles held by investors. The proportion wealth in mutual funds and money market funds changes from 74% and 20% pre-advice to 95% and 2% post-advice, respectively, effectively reducing cash positions. The increase in mutual fund holdings is also financed by reducing holdings in individual stocks and ETFs. The service reduced the small fraction of holdings in ETFs mainly because it was based on mutual funds during our study period.[3]

Our setting is very unique because, after signing up for the robo-advising, investors' portfolio changes are mechanical, that is, they do not involve investors' decisions. Conditional on signup, then, it is difficult to argue that the effects we document are not the direct result of the robo-advisor. Such

[3]As of 2019, the service invests clients in ETFs instead of indexed mutual funds.

argument would imply that investors would be knowledgeable enough to operate these large changes in their investment portfolios by themselves when they sign up for the robo-advisor, but not before signing up for robo-advising. It would also imply that—endogenously—individuals with similar characteristics would converge to the same portfolio allocations after signing up for robo-advice, irrespective of their portfolio allocations before signing up for robo-advice. In other words, the mechanical nature of the robo-advisor and the sheer magnitude of the portfolio changes are such that it is difficult to conceive that some unobservables factor is the ultimate driver of these portfolio changes.

We also estimate whether advice affects investors' investment performance. Because the service is relatively new, we can only analyze the near-term effects on portfolio returns. As a measure of performance, we use the annualized abnormal Sharpe Ratio, that is, the difference between the realized Sharpe ratio of each investor across all accounts and the realized Sharpe ratio of the market portfolio, where the latter is computed as the value-weighted returns on the NYSE/AMEX/NASDAQ CRSP portfolio. Irrespective of the horizon and the specification, we find investors' performance increases after adopting robo-advice. For example, at the 6-month horizon, the post-advice annualized Sharpe ratios average 0.115, statistically different from zero at the 1% level, while the pre-advice Sharpe ratios average -0.014, significantly different from zero at the 1% level, and the two are statistically different from each other at the 1% level.

We also find a positive and significant effect of signing up for advice and risk-adjusted performance in panel regressions that control for individual and time fixed effects. The improvement in performance is economically and statistically significant, starting from the second month after adoption of advice all the way to 36-months—the maximum horizon we can observe using our data. Furthermore, we find the improvement in performance is mainly driven by a reduction in investors' portfolio risk rather than an increase in investors' portfolio average returns. Finally, we find a strong and positive relation between robo-advising and risk-adjusted performance when we use an identification strategy that controls directly for the endogenous decision to sign up for advice.

The average results computed across all investors hide considerable cross-sectional heterogeneity. In an effort to understand which customers are more likely to benefit from robo-advising, we explore the cross-section of investors using a machine learning algorithm known as Boosted Regression Trees (BRT). BRTs allow us to analyze non-parametrically what investor characteristics are valuable in explaining the cross-sectional variation in the changes in portfolio allocations as well as the changes in investment performance pre- and post-advice.

For portfolio changes, the two most important investor characteristics influencing cross-sectional variation are the proportion of wealth held in equities by the investor at sign-up and the age of the investor. We find a very strong and negative relation between the change in the equity share and the fraction of wealth in equities at sign-up. Investors with no wealth in equities experience an increase in the share of equities of 30%. On the other hand, the advice service decreases by 30% the share of equities for those investors with 100% of their wealth in equities. The advice service systematically increases the equity exposure of investors who are less than 55 years of age and decreases the equity exposure of investors who are more than 55 years of age, confirming the stylized fact that older investors are systematically over-invested in equities while younger investors are under-invested in equities (e.g., see Guiso, Haliassos, and Jappelli, 2003). The change in equity exposure is economically large. It averages approximately -12% for the investors over 60 years old and almost +15% for those under 40 years of age.

We find that a large number of investor characteristics are related to the change in performance pre- and post-advice. Among them, we highlight that the cash share and the traded volume at sign-up are positively related to the improvement in performance post-advice, indicating that those investors who were trading substantially and/or were holding a very large portion of their wealth in cash, benefit more from advice. Other economically important relations are those associated with investor tenure as self-directed investors, the percentage of wealth in mutual funds, and the percentage of mutual fund holdings in index funds. In all of these cases, BRTs uncover a negative relation. Finally, the last and most relevant covariate is the share of equities held. The positive monotonic relation suggest that advice increases investors' performance more for those with higher equity shares, indicating that the service invests in a portfolio of mutual funds with higher risk-return trade-offs, compared to the average investor.

BRTs uncover in many cases strong non-linearities between regressand and covariates. To show that these are not the result of over-fitting, we perform an out-of-sample cross-validation exercise. BRTs do not overfit the training sample and they provide superior in- and out-of-sample performance, compared to linear models that use the same covariates. In fact, BRTs perform so much better than linear models in our setting that the out-of-sample performance of BRTs is superior to the in-sample performance of linear models.

In the third part of the paper, we move beyond investment performance and study investor attention before and after signing-up for advice. We show advised investors decrease the effort they need to exert to manage their investment portfolios. This reduction in attention is not related to an overall reduction in the investors' awareness of their financial condition, because investors tend to login more often to quickly acquire information regarding their portfolio wealth whenever they need to, but the overall time spent making investment decisions decreases after adopting advice.

We also study the determinants of adoption and attrition in robo-advice. The results suggest that it is the individuals who benefit the most from robo-advising—i.e. those who have low international diversification, have high expense ratios, and have high portfolio volatility as self-directed investors—that are the most likely to sign-up for advice and the least likely to quit the service. The results also suggest that the involvement of the human advisors increases dramatically the probability of sign up and decreases attrition.

1 Related Literature

Our work contributes to multiple strands of the finance and economics literatures. First, we contribute to the nascent literature in robo-advising. D'Acunto, Prabhala, and Rossi (Forthcoming) are the first to analyze the effects of robo-advising on the portfolios of individual investors. They find both promises and pitfalls, in that not all customers gain from adopting robo-advising. Our paper differs from D'Acunto, Prabhala, and Rossi (Forthcoming) in many respects. First, the robo-advisor in D'Acunto, Prabhala, and Rossi (Forthcoming) is a portfolio optimizer for stock portfolios. The robo-advisor analyzed here is of more modern conception. It uses indexed mutual funds rather than individuals stocks. A second major difference is that the robo-advisor analyzed here automatically trades for the investor, while the one in D'Acunto, Prabhala, and Rossi (Forthcoming) executes trades only when the investor logs-in into the platform and uses it. Reher and Sun (Forthcoming) study the effect of an Automated Financial Management service. They find that automated portfolios are more diversified that self-managed portfolios and that reducing the minimum balance required to gain access to the service increases customer fund inflows. Our results complement the ones in Reher and Sun (Forthcoming) as we assess the risk-adjusted investment performance and the type of portfolios and instruments held pre- and post-adoption of robo-advising. We also analyze the full cross-section of customers to measure which investors benefit the most.[4]

Bhattacharya et al. (2012) show individuals rarely follow unbiased—and beneficial—financial advice. They provide advice to a sample of German households and show very few households follow the advice provided. They provocatively conclude "*You can lead a horse to water, but you can't make him drink*." Our results stress that automatic implementation of advice is crucial for the efficacy of any form of financial advice, whether it is human- or robo-generated.

Second, and more broadly, we contribute to the household finance literature. Campbell (2006) argues financial markets are beneficial to households only to the extent that they participate in financial markets and they hold instruments that provide them with well-diversified investment portfolios. As shown in Badarinza, Campbell, and Ramadorai (2016), although households with higher socio-economic status conform more to theoretically optimal portfolio allocations, there are significant and persistent behavioral differences across countries. The field of cultural finance has related limited market participation to cultural norms and historical developments (D'Acunto, 2018a, D'Acunto, 2018b, and D'Acunto, Prokopczuk, and Weber, Forthcoming). Relatedly, Guiso, Sapienza, and Zingales (2008) study the role of trust in financial institutions and stock market participation. They show less trusting individuals are less likely to buy stocks and, conditional on buying stocks, they are likely to buy less of it. Our results show robo-advising can be a simple and inexpensive method to provide individuals with well-diversified portfolios and quickly increase exposure to domestic and international equities and fixed income securities.

Financial advising can potentially help mitigate under-diversification and help investors realize better outcomes (Gennaioli, Shleifer, and Vishny, 2015). However, for many retail investors traditional financial advisers are too costly. In addition, using data from the Canadian advice market, Linnainmaa et al. (2018) show the increased risk-taking on the part of the clients does not compensate for the higher costs associated with employing a financial adviser. Moreover, advisers often adopt a one-size-fits-all approach and might be prone to behavioral biases or display cognitive limitations (Linnainmaa, Melzer, and Previtero, Forthcoming). Fintech robo-advising gives clients access to financial advice at low cost. While robo-advising tools might be subject to the biases, conflicts, and limitations of the humans and institutions that develop them, they are by construction less influenced by the idiosyncrasies of specific

[4]For a comprehensive review of the robo-advising literature as well as the new developments related to the effects of robo-advising on financial inclusion, see D'Acunto and Rossi (Forthcoming) and Reher and Sokolinski (2020).

human advisers.

Our study is also relevant to the broader literature on technology adoption. Romer (1990) and Aghion and Howitt (1992) argue the adoption of new technologies are crucial determinants of economic growth. Comin and Mestieri (2014) argue there is a paucity of studies based on micro-data that measure the direct impact of technological progress on individuals' welfare. Our study contributes to this literature by providing new evidence in the context of financial advice.

2 Robo-Advising and Portfolio Characteristics

In this section, we first describe the data sources used in the study. We then present demographic and portfolio characteristics of robo-adviser investors before they sign-up for the advice service. Third, we show how the portfolio characteristics of advised investors change over time after they sign up for the robo-advising service. Finally, we analyze the type of assets advised and non-advised investors are invested in. All results are computed at the investor level and include all account types, that is, taxable and non-taxable (IRA) accounts.

2.1 Data

The study uses anonymized proprietary data from Vanguard. The data contains information on trades, positions, demographic characteristics and investor-advisor mapping information for previously self-directed investors that have interacted with Vanguard's Personal Advisor Services (PAS).

The main results are computed using the sample previously self-directed investors who signed up for advice between January 2015 through December 2017. At the time of sign-up, we have a total of 82,526 investors. The number of investors with at least 12 months of data as advised investors is 54,747.

The trades data includes all the trades placed over the period January 2015 through December 2017. The positions data contains monthly holdings observations for the same investors. The demographic characteristic data contains information on investor marital status, age and gender, together with detailed information regarding the dates on which the investor initiated, enrolled, implemented and quit the advice service. The investor-advisor mapping data contains information on the interactions—

i.e. meetings, phone calls, etc.—between investors and human advisors.

The study also uses a variety of additional data sources. Stock market information such as prices, returns and trading volumes – among others – is obtained from CRSP, and CRSP Mutual Funds. In addition, the CRSP Mutual Funds database contains information regarding mutual fund fees, turnover, expense ratios, investment allocations, degree of indexation and the mutual fund classification provided by Lipper.

2.2 Demographic and Portfolio Characteristics Pre-Advice

We start by reporting demographic and portfolio characteristics of the investors that sign up for advice, computed the month before the investors sign up for the service. The results are reported in Table 1, where for every variable we report mean, standard deviation and various percentiles of the distribution—ranging from the 10^{th} to the 90^{th} percentile. Panel A focuses on the demographic characteristics. The average investor is 63 years old and the median is 65; 53% of the users are males and 69% of them are married. Tenure as self-directed investors varies a lot. It ranges from two years at the 10^{th} percentile to 26 years at the 90^{th} percentile. For comparison, the average investor age is 51 in Gargano and Rossi (2017) and Barber and Odean (2001). The percentage of women, which equals 46%, is larger compared to both Gargano and Rossi (2017), 27%, and Barber and Odean (2001), 21%. At approximately 14 years, average investor tenure is also longer, compared to other brokerage account datasets in the literature. Average investor tenure in Gargano and Rossi (2017) is less than 9 years.

Panel B of Table 1 reports results for portfolio allocation. Investors' wealth is substantial. It averages \$588,246 and is heavily skewed to the right. The median invested wealth is \$282,450. The number of assets per investor is 7.8 and the median is 5. It may appear that these investors are substantially under-diversified, but this is really not the case, because many of these investors are very heavily invested in mutual funds. On average, 72% of the wealth is invested in mutual funds (and another 20% in money market funds) rather than individual stocks, so investors are likely to be very diversified, even if they hold only 5 assets.

The average investor has 54% of his/her portfolio invested in equities, followed by 24% in bonds, and 22% in cash—mainly money market mutual funds. These averages hide a very large cross-sectional variation, with almost 10% of the investors almost completely invested in equities and 15% of the investors invested only in bonds and/or cash. Stocks and bonds are not held directly, but mainly through mutual funds. In fact, 72% of the wealth is invested in mutual funds, followed by cash at 20%. Only 3% of investors' wealth is held in individual stocks and 3% in ETFs. Finally, only a negligible number of investors have direct exposure to corporate bonds and options (not reported in the table).

Mutual fund holdings can be decomposed according to the fund strategies. As reported at the bottom of Panel B, 47% of mutual fund holdings are in indexed mutual funds, while 10% of mutual fund holdings are in funds that invest internationally. Only a negligible number of customers invest in mutual funds with standalone emerging markets exposure.[5]

Panel C focuses on fees and transactions. Starting from mutual fund fees, the average management fee is 14 basis point, but some investors spend as much as 58 basis points a year in management fees. The expense ratio results are similar. The average is 0.19, the median is 0.14, and some investors have expense ratios close to 1% per year. The third row of Panel C focuses on the turnover ratio of the mutual funds held, that averages 0.32. In terms of active transactions, investors place on average 3 transactions per month, for an average of $85,000 dollars.

2.3 Demographic and Portfolio Characteristics Post-Advice

We now report how investors' portfolio allocations change after signing up for the advice service. In Table 2, we compute the same quantities of Table 1, but focus on the 12 months after adoption of advice. The demographic characteristics such as age, tenure, proportion of males and married people are reported in Panel A of Table 2 and are—overall—rather similar to the ones in Panel A of Table 1. Panel B reports the portfolio allocation results. At $758,000, average wealth is higher than in Table 1. This is the result of stock market appreciation and investors' contribution to their portfolio. The number of assets in each account increases slightly from 7.79 to 7.95. The percentile distribution shows that advice shrinks the number of stocks held in the tails of the distribution. The 90-th percentile of the number of assets held in each account drops from 17 in Table 1 to 14 in Table 2.

Continuing with the results in Panel B, portfolio allocation is where we observe strong changes, particularly in the allocation to bonds and money market mutual funds (cash). The percentage of bonds

[5] We isolate indexed mutual funds using the "IndexFlag" from the CRSP mutual fund database. We also identify the funds with international exposure as the ones classified as either "international" or "global" by the Lipper classification. Finally, we identify the emerging markets funds using the "emerging" Lipper classification.

increases by 15 percentage points to 39%, while the allocation to cash decreases by 19 percentage points to only 2%. Finally, the equity share increases by 5 percentage points to 58%. The next four lines in Panel B of Table 2 focus on the investment vehicles used. Almost all of investors' wealth—95% of it—is invested in mutual funds, with almost no share of wealth in money market mutual funds (2%), ETFs (1%), or individual stocks (1%).

Advice has a very large effect on indexation and international diversification as well. Before advice, the average investor has 47% of their wealth in index funds. This increases to 83% after signing up for advice. We find a similar effect for investor's exposure to international markets, that increases from 10% to 32%. Interestingly, we do not find much of an effect in terms of emerging markets exposure that, if anything, declines after advice. As we show in Online Appendix A, this is because international mutual funds (VTIAX, for example) have emerging market exposure that is not well captured by the Lipper classification.

Panel C of Table 2 shows advice moves investors to passive mutual funds, characterized by lower fees and turnover ratios. Management fees are halved, from 14 to 7 basis points, while the expense ratio is reduced by more than 50% as it drops from 19 to 9 basis points. The turnover ratio instead drops by approximately 20%, from 0.32 to 0.27.

The results in Tables 1 and 2 are computed using all accounts available either at robo-advice adoption (a total of 82,526 clients) and at the 12-month after robo-advice implementation (a total of 54,747 clients). This procedure has the advantage of maximizing the number of observations. However, the results mix the effect of advice with the sample composition effect—we do not have as many investors 12 months after signing up for advice compared to the month before signing up.

In Figures 1 and 2, we show the time-series behavior of the most important quantities before and after signing up for the service. In each plot, the blue line represents average values while the red dashed lines are 95% confidence intervals for the mean. Time "0" represent the month before investors sign up for advice. Figures 1 and 2 have the advantage of purging the results from any composition effect, as we only track investors that survive over the 24-month window of PAS adoption. The drawback is that the number of investors used to compute the results is smaller.

Subfigures (a), (b), (c), and (d) of Figure 1 show the time-series behavior of the changes for bond, cash (including money market mutual funds), equity and mutual fund holdings. We highlight several findings. First, it takes 1-2 months for the service to converge to the new portfolio allocations post adoption. Second, in the months leading to signing up for advice, investors tend to change their investment portfolio. This is most evident in Panels (b) and (d), that is, for the results related to the percentage of cash held and the percentage of mutual funds held. The first one trends up while the second trends down. This has—potentially—important implications when it comes to evaluating investors' performance and characteristics pre- and post-advice. Second, the results in Figure 1 are very much in line with the ones in Tables 1 and 2.

Figure 2 presents results for indexation, international diversification, expense ratio and trading volume. In all cases, the changes take place over the course of 1 or 2 months and are in line with the results in Tables 1 and 2. For the indexation results, Subfigure (a), we observe some drifting, starting 6 months before the adoption of advice. The trading volume results are unique as they display marked non-monotonicities. Trading volume spikes for approximately 1 to 2 months after enrollment into the service as the advice service changes investors' positions to the new target weights.

Figures Online I and Online II repeat the computations in Figures 1 and 2, but focus on the median and the 10th and 90th percentiles. The results show a very large reduction in the cross-sectional variation across investors post-adoption of advice. The reduction is, in certain instances, rather dramatic. If we focus on the percentage of equities in the investor portfolios (Subfigure (c) of Figure Online I), the 10th to 90th percentile of the distribution ranges from 0.25 to 1.00 twelve months before adoption to 0.40 to 0.85 after adoption. Also very large is the reduction in the cross-sectional dispersion across investors when it comes to indexation, percentage of wealth in international funds, and expense ratios, reported in panels (a) through (c) of Figure Online II.

In Section Online Appendix A.1 we present a detailed analysis of the actual tickers held by advised and non-advised investors. Advised investors have portfolio allocations that are more homogenous—they hold similar tickers. Furthermore, a large part of advised investors' wealth is placed in a few (low-cost) indexed mutual funds that focus on US Equities (28% of total wealth), International Equities (18% of total wealth), US Bonds (15% of total wealth) and International Bonds (11% of total wealth).

3 Performance before and after Advice

The portfolio allocation results reported so far suggest that the advice service may improve investors' performance as it places account-holders in diversified US and international low-fee indexed mutual funds. It also reduces investors' cash holdings. In this section, we provide a comprehensive analysis of the pre- and post-advice investment performance.

3.1 Pre- and Post-Advice Statistics

Because the service is relatively new, we can only analyze the near-term effects on portfolio returns. As a measure of performance, we start by using the annualized abnormal Sharpe Ratio, that is, the difference between the realized Sharpe ratio of each investor across all accounts and the realized Sharpe ratio of the market portfolio, where the latter is computed as the value-weighted returns on the NYSE/AMEX/NASDAQ CRSP portfolio. For each account-holder, portfolio returns and volatilities are computed using beginning-of-month investment holdings. Furthermore, portfolio volatilities are computed as realized volatilities using squared daily returns. Performance is computed starting at adoption in Table 3 and starting 6 months before and after advice is implemented in each account in Table 4. The latter is to make sure that users' portfolio allocations are in steady-state before and after their sign-up, as we discuss in Section 2.3.

The results, computed across all investors, are reported in Table 3. Panel A reports results for annualized Sharpe ratios computed at the 3-month horizon. Panels B and C reports results computed at the 6-, and 9-month horizons. Each column within each panel reports average abnormal Sharpe ratios, t-statistics testing the null hypothesis that the averages are equal to zero, and the number of observations used in the computations of the results. The first and second column of each panel report average abnormal Sharpe ratios after (column 1) and before (column 2) signing up for advice, computed across all investors available. Columns 3 and 4 repeat the the exercise only for those accounts in the sample both before and after—we refer to them as "matched" investors. The last column reports the average performance difference after and before signing up for advice for the matched investors. In all cases, we subtract advice fees from the investors' performance.[6]

[6]Fees for the robo-advising service we study start at 30 basis points per year and decrease as a function of assets under management. To be conservative, we deduct 30 basis points a year from all accounts—which understates the performance of wealthier individuals.

Across all panels we find the performance improvement is sparse. This should not be surprising as it takes a couple of months for the advice service to reach the new portfolio allocations.

The corresponding results starting the computations 6 months before and after advice adoption—Table 4— are instead much stronger, as expected. At the 3-, 6-, and 9- month horizons, we consistently find that the abnormal Sharpe ratio is positive and significant, with values ranging between 0.094 and 0.115, when computed across all investors. The pre-advice Sharpe ratios are instead negative and significant at the 3- and 6-month horizon and positive, but economically small, at the 9-month horizon. The matched results show similar and consistent results: the post-advice performance is superior than the pre-advice performance, with the highest value equaling 0.432 at the 9-month horizons, statistically significant at the 1% level.

In Figure 3, we focus on the results reported in the last column of Panel C of Table 4 and report the full distribution of realized abnormal Sharpe ratio differences before and after adopting advice. The density has a mean of 0.323 (blue solid line), which is statistically different from zero (red dotted line) at the 1% level. Figure 3 conveys two facts. First, it is not only a handful of investors that improve the performance post-advice, but the majority of investors have a higher realized abnormal Sharpe ratio difference post-advice. Second, there is a large dispersion in performance gains. These are in part driven by idiosyncratic reasons—some investors may have held securities that perform very poorly or very well before signing up for advice and in part driven by the systematic suboptimal portfolio allocations on the part of the investors. We explore what investor characteristics pre-advice relate to greater and smaller performance improvement post-advice in Section 4.2.

3.2 Robo-Advising on Investor Performance: Panel Regression Estimates

The results in Section 3.1 control for time-variations in stock market returns and risk, because we construct abnormal Sharpe ratios. However, they do not control for investor characteristics. They also do not decompose the effect of robo-advising on portfolio risk and portfolio returns. We overcome these limitations by resorting to panel regressions that include time-effects and individual fixed-effects. In particular, we estimate the baseline specification:

$$Sharpe_{i,t} = \alpha_i + \beta_t + \gamma \ ROBO_{i,t} + \epsilon_{i,t}, \tag{1}$$

where $Sharpe_{i,t}$ is the realized Sharpe ratio over month t, α_i denote investors' fixed effects, β_t are time-effects, and $ROBO_{i,t}$ is a dummy variable that equals 1 when the investor's portfolio is managed through the robo adviser and zero when the investor is self-directed. The coefficient of interest γ captures the effect of advice on the risk-return of each investor's portfolio, controlling for investors' characteristics and time-variation in the risk-return trade-off realized in the market. The results are reported in Table 5. Panel A reports results for investors' Sharpe ratios. Specifications 1 and 2 do not include individual fixed effects, while specifications 3 through 6 do. Specifications 1 through 4 do not include time-effects, while specifications 5 and 6 do. Finally, in specifications 1, 3, and 5, standard errors are clustered at the individual level. Standard errors are instead double-clustered by both individual and time in specifications 2, 4, and 6.

We note three facts. First, it is important to control for both investor and time-effects. The coefficients change from 0.88 to 2.05, with the introduction of individual fixed-effects. They further change from 2.05 to 0.364 once we introduce time effects. Second, double-clustering the standard errors is crucial to account for the correlation among different investors in the same month and different months for the same investor. Test statistics are 30 to 150 times larger if we do not double-cluster the standard errors. Third, and most importantly, the coefficient 0.36—in preferred Specification 6—is economically rather large, given that the realized market Sharpe Ratio over the period January 2015 to December 2017 was 0.88.

Panel B and C repeat the exercise, but focus on portfolio returns and volatility, respectively. The results in Panel B show there is virtually no relation between portfolio returns and robo-advising adoption. The results in Panel C show instead a very large and significant relation between the reduction in portfolio risk and robo-advising. Over the period January 2015 to December 2017, the annualized monthly stock market volatility averaged 12%. A coefficient of -0.009 implies a volatility reduction of $-0.009/0.121 = -0.074$, or 7.4%.

The results reported in Table 5 include a dummy variable that equals one for all the periods after the adoption of the robo-advisor. To show how the effect of the robo-advisor evolves since its adoption, we estimate the following alternative to Equation (1):

$$Sharpe_{i,t} = \alpha_i + \beta_t + \sum_{j=-5}^{35} \gamma_j \ ROBO_{i,j,t} + \epsilon_{i,t}, \tag{2}$$

where all quantities are defined as in Equation 1, except that the dummy variable $ROBO_{i,j,t}$ is equal to zero for every month, except for the j-th month before and after adoption. The $0th$ month is the one when the robo-advisor is adopted, negative values of j refer to the months before advice is adopted, and positive values of j refer to months after robo-advice is adopted. We start 5 months before robo-advising is adopted to show that there is no systematic under-performance or over-performance on the part of the investors that adopt robo-advising. We stop at the $35th$ month post-adoption, as it is the longest horizon for which we have robo-advised clients. Rather than reporting the γ_j's coefficients and standard errors in a lengthy table, we report coefficients and 95% confidence intervals based on double-clustered standard errors in Figure 4.

Three facts are worth noting. First, dummy coefficients show a significant improvement in performance after adoption. All the coefficients after the $0th$ month are significantly different from zero, except for the $34th$ month coefficient, that is significant at the 10% level, but not at the 5% level. This is probably due to the low number of investors for which we have close to three years of investment portfolio data. Second, the results for the $0th$ month indicate no-significance outperformance in the transitional month. This is because the transitional month entails a lot of trading to achieve the new portfolio allocations. Third, none of the pre-advice dummies are significant, suggesting that those adopting robo-advice did not significantly under- or over-perform pre-adoption in the months before adoption. Because the pre-adoption dummies may not fully control for investors under- and over-performance pre-adoption, we report next the results of a formal identification strategy that controls for the endogenous timing of investors sign-up. The results confirm the main findings reported here.

3.3 Identification Strategy Results

The results presented in Section 3.2 do not control for the fact that individuals endogenously decide when to sign-up for advice. Because the decision is endogenous, it may be that individuals sign up right after performing poorly in ways that are not fully controlled for by the pre-adoption dummies in Specification 1. This selection would imply that the benefits we uncover in Section 3.2 in terms of risk-return trade-off are overstated.

We tackle this concern using an identification strategy that exploits two main features of the robo-advisor under consideration. First, the portfolio holdings of the investors are rather stable over time, because the robo-advisor invests mainly in indexed mutual funds. Second, while the robo-advisor allows for portfolio personalizations, it also standardizes investors' portfolio allocations post adoption. Both these features are exploited in the strategy described below.

We start by considering only those individuals who have signed-up for advice during the sample and focus on their returns starting from 6 months after they have signed-up for advice. We then construct, for each investor, the correlation of the monthly realized sharpe ratio with every other investor who is advised, using all the available months of overlap to conduct the calculations. For each investor, we rank all other investors in terms of their risk-return correlations when they are advised. Intuitively, investors that have the same risk-return profiles when advised are very similar in their characteristics—at least from the robo-advisor perspective. Finally, we exploit the fact that investors sign up at different times to construct counter-factual returns for each investors. We provide a simple example of the strategy below.

As shown in the diagram below, suppose we have three investors A, B and C and we want to construct the counterfactual performance for investor A. Suppose also that A signed up for advice on January 1st 2016, while B and C signed up for advice on July 1st 2014. Finally, assume that the three investors are advised through the end of 2019. We use the period starting from July 1st 2016 through December 2019—the are highlighted in green—to compute the realized Sharpe Ratio correlations of investor A with both investors B and C. Suppose that the correlation with B is 92% while the correlation with C is 99%. For investor A, we would rank investor C as first and investor B as second in the construction of the counterfactual. In the final step, we would compare the performance of investor A during the period January 1st 2015 through December 31st 2015—when he/she was unadvised—to the performance of investor C over the same period—when he/she was already advised. In the diagram below, this is denoted using the red shaded area.

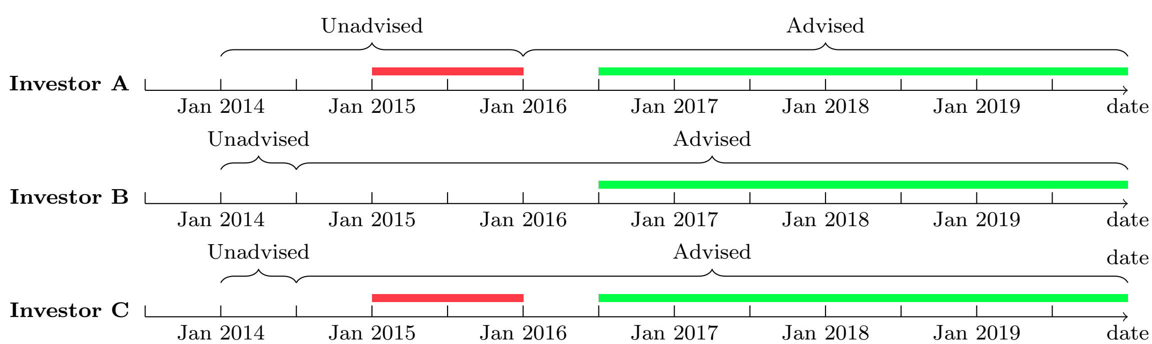

In the baseline specification, we only use the closest matches for each investor—the decile of investors with the closest portfolio realized Sharpe ratio. In alternative specifications, we allow for more investors in the control group. At the other extreme, we include the average performance of all the advised investors as the control group. The trade-off is that, as we include more and more investors in the control group, the risk-return performance of the control group become more stable, but potentially less representative.

This strategy allows us to construct counter-factual returns for investors for periods where they did not endogenously decide to sign-up, attenuating the concern that the endogenous decision to sign up for advice biases our baseline results. The results are reported in Table 6.

The first column compares the performance of the unadvised individuals to the top decile of closest matches. The second and third column focuses on the top 2 and 3 deciles, respectively. Finally, the fourth column focuses on all matched investors. In all cases, we find a significantly negative difference between the performance of unadvised and advised clients.

Overall the results of our identification strategy are in line with those reported in Table 5 and indicate that endogenous sign-up decisions are unlikely to be the main drivers of our results.

4 Using Machine Learning to Assess the Effects of Robo-Advising

The previous sections analyzed how the advice service changed the investment portfolios and the investment performance of the average investors. However, we can expect the robo-adviser to have a differential impact on the portfolio allocation and the performance of each investor, depending on investor characteristics at sign-up.

As a motivating example, we work with the share of equities across investors. In Figure 1, we showed the average change in the equity share was rather small across investors. This result, however, hides a very large heterogeneity across investors. To illustrate the point, we report in Figure 5 portfolio changes for investors with low (less than 10%) and high (more than 90%) equity shares before signing up for the service. In both cases, advice resulted in a major portfolio overhaul. In the first case, advice increased equity holdings from approximately 5% to almost 50%. In the second, it decreased it from 95% to approximately 70%.

The portfolio changes operated by the robo-adviser are largely a function of the investment portfolio of the investors at sign-up as well as investor preferences and demographic characteristics. For example, older individuals are likely to be assigned a lower share of their wealth to risky assets, while younger individuals a higher share. Investors' lifestyle may also play a role: namely, investors with different projected expenses relative to wealth are likely to be assigned different investment portfolios. Finally, investors' preferences such as risk aversion play a role. The final portfolio allocation of each investor is the product of investors' characteristics, the algorithm used by the robo-advisor and the customizations operated by the human advisor upon investors' request. It is therefore difficult to know what factors ultimately play a role.

A standard way to analyze this problem would be to use linear regression, but it is not clear that investors' demographic and portfolio characteristics are linearly related to the changes in investors' portfolios. It is also not clear *ex ante* what factors would be relevant. The result of running a kitchen-sink regression is that we would likely run the risk of overfitting the data and estimate spurious relations between regressors and regressand. Instead, we use a machine learning method known as Boosted Regression Trees. Boosted Regression Trees not only allows large conditioning information sets, but it also allows for non-linearities—all without overfitting or falling prey of the so-called curse of dimensionality. We provide a brief introduction of BRT—including an introduction to the concepts of *partial dependence plots* and *relative influence measures*—in Online Appendix B.1.

4.1 Which Investors Experience the Biggest Portfolio Changes?

This section uses BRTs to decompose the effect of advice on investors portfolio allocations. As a measure of portfolio change, we adopt the share of equity held by the investor. Intuitively, from Figure 5 we know that investors with a very high equity share before signing up for advice are likely to experience a decrease in their equity share. Those with a low equity share are instead likely to experience an increase in their equity share.

In order to decompose and visualize how advice changes investors' portfolio allocations after sign-up, we estimate a BRT model with 10,000 boosting iterations. The dependent variable is the change in the share of equities six months before and after signing up for advice. As conditioning variables we use a total of 15 regressors, divided into three groups. The first group contains demographic characteristics: Age, the age of the investor as of December 2017; $Male$, whether the investor is male; $Married$, whether he/she is married; $Tenure$, the tenure of the investor as of December 2017. The second contains regressors related to portfolio characteristics: $NumAssets$, the number of assets held by the investor across accounts; $PctEquityShare$, the percentage of wealth in equities—held directly or through mutual funds; $PctCashShare$, is the percentage of wealth money market mutual funds—held directly or through mutual funds; $PctMutualFunds$, the percentage of wealth directly invested in mutual funds; $PctStocks$, the percentage of wealth directly invested in individual stocks; $PctETF$, the percentage of wealth directly invested in ETFs; $PctIndex$, the percentage of mutual fund wealth invested in index funds; $PctEmerging$, the percentage of mutual fund wealth invested in emerging market funds—identified using the Lipper mutual fund classification. The third groups relates to variables related to transactions and fees paid: $MgtFees$ are the value-weighted management fees charged by the mutual funds held by the account-holders; $Transaction$, the number of transactions directly initiated by the investors over the month before signing up for advice; $Volume$, the volume (in US dollars) traded by the investors over the month before signing up for advice.

Out of the 15 predictor variables, only three variables have a relative influence higher than 1%. The variable $PctEquityShare$ has the highest relative influence measure, totaling 81.9%. This means that the splits based on $PctEquityShare$ contribute to 81.9% of the reduction in the empirical error of the model. The second variable is Age, that has a relative influence of 15.6%. Finally, the third variable is $PctCashShare$ that has a relative influence of 2.1%. The remaining covariates explain very little of the variation in investors' equity share post-advice. This indicates that the change in risky share is mainly determined by the investors' positions when signing up for the advice service and investors' age.

We report the univariate partial dependence plots for the three most important predictors in Panel A of Figure 6. The first subfigure reports results for the share of wealth in Equities pre-advice. There is clearly a negative relation between the change in the equity share and the fraction of wealth in equities pre-advice. Those investors with no wealth in equities experience an increase in the share of equities of 30%. On the other hand, advice decreases by 30% the share of equities for those investors with 100% of their wealth in equities. Interestingly, we find that that the partial dependence crosses the "0" on the y-axis, when the Equity Share equals approximately 0.6, indicating that advice does not change the equity positions for those investors that already have roughly a 60-40 split between equities and bonds.

The second subfigure plots the partial dependence with respect to Age. Systematically, advice increases the equity exposure of investors that are less than 55 years of age and decreases the equity exposure of those investors that have more than 55 years of age. The reduction in equity exposure is significant. It equals approximately -12% for the investors over 60 years old and almost +15% for those under 40 years of age. The third covariate is Cash share. The relation is once again positive, but economically small.

Panel B of Figure 6 present the bivariate dependence plots. The first plots the partial dependence of the change in equity share with respect to both the pre-advice equity and cash share. The bivariate plot shows the joint negative relation between both regressors and the changes in the Equity share. It also shows that the relation between age and the change in equity share is monotonic, but not linear. The second plot instead displays the partial dependence plots of the change in equity share with respect to the equity share and the cash share. The plot that that, jointly, the change in equity share is negatively related to pre-advice equity share and positively related to the cash share.

One interesting finding of the partial dependence plots is that the both the bond and cash share are linearly related to the changes in equity share operated by the robo-adviser. The only regressor that instead displays a monotonic, but not linear, relation is Age. As a result, it could well be that a linear regression could work almost as well as BRT in this case. Even in this rather simple setup, however, linear regression results are misleading. A kitchen-sink linear regression model estimates as significant at the 5% level seven of the fifteen regressors: Age, $Male$, $Married$, $NumAssets$, $PctEquityShare$, $PctCashShare$, $MgtFees$. Some of these regressors, such as $Male$ and $Married$, are likely significant because they compensate for the nonlinearities in the variable Age. We know it is the case, because they become insignificant when we include higher-order transformations of the variable Age in the linear model. If an econometrician were to estimate a kitchen-sink regression and did not have access to the results of BRT, however, he/she would conclude that factors such as gender and marital status have an impact on the portfolio changes the robo-advisor operates, even though this inference would be only driven by the mis-specification of the linear model.

The change in portfolio share is a rather easy quantity to model as it is likely to be a deterministic—even though unknown—function of investors' demographic and portfolio characteristics at the time the investors sign up, as well as their risk-preferences, liquidity needs, and employment characteristics, which we cannot observe. A more challenging question is whether we can use investors' observable characteristics to predict which users are likely to benefit the most from robo-advising. We undertake this analysis next.

4.2 Which investors Benefit the Most from Robo-advising?

In this section, we explore whether we can explain the cross-section of changes in risk-adjusted performance pre- and post-advice using investors' characteristics at the time of sign-up.[7] We use the same setup and covariates as in Section 4.1, the only difference being that we replace the share of bonds with the share of equities as regressor, to ease the interpretation of the results.

The dependent variable is the change in the abnormal Sharpe ratio six months before and after signing up for advice, each computed using 9 months of data—see Figure 3 for details. The task in this section is much more challenging than the one in the previous section, for at least two reasons. First, abnormal Sharpe ratios are noisier, as they are computed on the basis of only 9 months of returns and realized volatilities. Second, the change in abnormal Sharpe ratios is likely to not only be driven by the equity-bond allocation decision, but also by the characteristics of the individual securities held. Intuitively, in this section we explain the cross-sectional variation in the performance improvement post-adoption reported in Figure 3.

The relative influence results are reported in Figure 7, while the partial dependence plots for the top 9 covariates by relative influence are reported in Figure 8. Among the top 9 covariates, some of the relations are immediately intuitive. For example, the cash share (relative influence of 11%) and the traded volume (relative influence of 3.4%) at sign-up are positively related to the improvement in performance post-advice. This is simply saying that those investors that were trading a lot and/or where holding a very large portion of their wealth in cash, benefit more from advice.

Other economically intuitive partial dependence relations are those associated with investors' tenure as self-directed investors (relative influence of 13%), the percentage of wealth in mutual funds (relative influence of 5.4%), the percentage of mutual fund holdings in index funds (relative influence of 2.2%).

[7] As noted in the introduction, we do not focus on other sources of value investors may derive from financial advice, such as financial education/literacy benefits. We also do not account for emotional/hedonic rewards such as improved financial well-being and peace of mind.

In all cases, BRTs uncover a negative partial relation. This suggests that those investors that were not holding a lot of indexed funds and were not holding a lot of their wealth in mutual funds, are the ones benefitting more from signing up for advice. The partial dependence with respect to investors' tenure suggest instead that less experienced individuals are the ones that benefit the most from advice.

BRTs uncover also markedly non-linear and non-monotonic relations between the change in risk-adjusted performance and investors' age (relative influence of 16.5%), the mutual funds' management fees (13.3%) and the number of assets held (relative influence of 3.6%). For the first two regressors, the relation is U-shaped. Investors benefitting the most are the ones in their forties and mid-fifties, and the very senior citizens, while there is a negative relation between age and change in post-advice performance for investors in the second half of their fifties and their sixties. The relation is probably due to the fact that advice tends to increase the equity exposure of the investors in their forties and mid-fifties, decrease it for investors in the second half of their fifties and their sixties, and leave them unchanged for the investors in their seventies (see the second Subfigure in Panel A of Figure 6).

The relation between management fees and change in performance is also U-Shaped, indicating that those customers investing in very expensive active funds as well as the few very cheap funds are the ones that benefit the most. The result is probably driven by the fact that over our sample period, actively managed funds have performed relatively poorly, compared to passive funds.

The relation between number of assets and performance change has instead an inverse U-shaped relation. This is due to the fact that individuals with few assets are likely to be holding mutual funds. Associated with a higher number of assets are instead those investors that invest in individual equities. These customers do benefit from advice is it increases their diversification. Those individuals that instead had 25 or 30 assets where likely to be already rather diversified, because they have a wide range of both stocks and mutual funds. They therefore do not benefit much from adopting the robo-adviser.

Finally, the last covariate and most relevant covariate is the share of equities held (relative influence of 27%). The positive monotonic relation suggest that advice increases investors' performance more for those with higher equity shares, indicating that the advisory service invests in a portfolio of mutual funds with higher risk-adjusted profiles, compared to the average investor.

4.3 In- and Out-of-sample Performance of BRTs

One of the main criticisms against non-parametric models is that they tend to overfit the training dataset. One could be worried that the non-linearities and non-monotonicities uncovered in Figure 8 and described in Section 4.2 are the result of BRTs fitting noise rather than the structural relation between the covariates and the dependent variable. We show here that this is not the case. Crucially, we also show that the most important free parameter, i.e., the number of boosting iterations, does not significantly affect the out-of-sample performance of the method.

To asses whether BRTs are overfitting the training dataset, we perform the following cross-validation analysis. We take the original dataset on which we estimate our BRT results and exclude half of the observations. We then estimate the BRT model on one half of the data and test its performance on the other half of the data. We repeat the analysis 100 times. On every iteration, we store the in- and out-of-sample performance of BRTs for boosting iterations that range from 100 to 20,000. For every iteration, we also store the in- and out-of-sample performance of a linear model that uses the same regressors as BRT. Finally, we report two sets of results. The first set reports the in- and out-of-sample performance of BRTs —averaged across all cross-validation rounds— for different boosting iterations. The second plots the density of the out-of-sample performance of the BRT model and the linear model across all cross-validation rounds.

To show how the performance changes depending on the setting, we report the results for the change in the equity share before and after signing up for advice in Figure 9 and the change in investment performance before and after signing up for advice in Figure 10

4.3.1 Changes in Portfolio Allocation. As mentioned in Section 4.1, the out-of-sample prediction of the changes in portfolio allocation is likely to be not very challenging, because the change in the risky share is likely to be some deterministic function of investor demographic and portfolio characteristics (which we observe) as well as well as investor preferences and tolerance of risk (which we do not observe).

As highlighted in Section 4.1, the partial dependence plots show that the relations between the regressors and the independent variable are mostly linear, with the exception of *Age* that appears to be monotonic, but not linear. As a result, we should expect the linear model to perform rather well compared to BRTs. This is indeed what we find.

Subfigure A of Figure 9 show the average in- and out-of-sample performance of BRTs and the linear model for portfolio changes. As the number of boosting iterations increases, the fit of BRTs improves and rises to almost 60%, as shown by the black line. For comparison, note that the linear model has an in-sample R^2 of only 57.2%—green line. The out-of-sample performance BRT improves as the number of boosting iterations raises from 100 all the way to approximately 12,000, as shown by the red line. The out-of-sample fit then asymptotes and stabilizes at around 57.25%, a value greater than the in-sample R^2 of the linear model. The out-of-sample fit of the linear model is instead worse, equalling 56.7%—blue line.

Subfigure A in Figure 9 reports averages across simulation rounds. To show how the out-of-sample performance of BRTs and the linear model compare, we report in Subfigure B of Figure 9 the density of the out-of-sample R^2 across simulation rounds. Consistent with the findings in Subfigure A, BRTs consistently outperform the linear model out-of-sample.

4.3.2 Changes in Investment Performance.

Explaining the changes in investment performance is likely to be rather challenging, because the realized Sharpe ratios are estimated using only 9 months of daily data and are therefore rather noisy. Also, over such short period of time, it is possible that certain stocks or portfolios will deliver very low or large returns for idiosyncratic reasons.

As shown in Section 4.2, the partial dependence plots show that the relation between the regressors and the independent variable is certainly non-linear and in some cases also strongly non-monotonic. As a result, we should expect the linear model to perform rather poorly compared to BRTs. This is indeed what we find.

Subfigure A of Figure 10 show the in- and out-of-sample performance of BRTs and the linear model. As the number of boosting iterations increases, the fit of BRTs improves and rises to almost 5.91%, as shown by the black line. For comparison, note that the linear model has an in-sample R^2 of only 1.70%—green line. The out-of-sample performance BRT improves as the number of boosting iterations raises from 100 all the way to approximately 10,000, as shown by the red line. The out-of-sample fit then asymptotes and stabilizes at around 1.92%, a value greater than the in-sample R^2 of the linear model. The out-of-sample fit of the linear model is instead rather poor, equalling 0.06%. If we allow for non-linearities in the linear model by including second- and third-order higher transformations of the original regressors, we find that the out-of-sample R^2 deteriorates even further to 0.05% and 0.04%, respectively.

In Subfigure B of Figure 10 we present the density of the out-of-sample R^2 across simulation rounds. Consistent with the findings in Subfigure A, BRTs consistently outperform the linear model out-of-sample.

5 Beyond Performance: Robo-advising and Investor Attention

So far, the paper has focused on the investment performance of investors before and after signing up for advice. We now move beyond investment performance and focus on the time spent by investors in managing their finances. The optimal inattention models of Abel, Eberly, and Panageas (2007, 2013) predict investors should pay attention to their investment portfolios to equate the cost of paying attention to the portfolio, i.e., time and cognitive costs, to the benefit of knowing what is the state of their investment portfolios. After adopting portfolio advice, investors are unlikely to need to login to their investment portfolio to operate any changes in their investment portfolio. The only reason they should be logging into their account is to monitor the value of their investment portfolio and possibly to monitor how the robo-advisor operates.

In this section, we focus on investors' logins and the time spent on the website. We estimate how investors' attention changes after signing up for advice, compared to before. We estimate the following baseline panel regression:

$$Attention_{i,t} = \alpha_i + \beta_t + \sum_{j=-5}^{35} \gamma_j \ ROBO_{i,j,t} + \delta \ X_{i,t} + \epsilon_{i,t}, \tag{3}$$

where all quantities are defined as in Equation 2, except that the dummy variable $ROBO_{i,j,t}$ is equal to zero for every month, except for the jth month before advice enrollment is initiated and after advice is implemented. The $0th$ month is the one when advice is implemented, and negative values of j refer to the months before the process of enrollment into advice is initiated. We start 5 months before the process of enrollment into advice is initiated to show there is no systematic under- or over-attention before signing up for advice. We stop at the $35th$ month post-adoption as it is the longest horizon for which we have advised clients. Note that in these results we do not include dummies for the periods of advice enrollment as the enrollment process ranges between one to several months across individuals. The regressor $X_{i,t}$ contains investors' portfolio return and risk, which do not have an economically significant effect on attention. The coefficients of interest are the γ's. Rather than reporting the γ_j's coefficients and standard errors in a lengthy table, we report coefficients and 95% confidence intervals based on double-clustered standard errors in Figure 11.

Figure 11 uses the days with logins within a month as a measure of attention. Panel A aggregates logins across all devices, while Panels (b) through (d) break down the results isolating attention through a desktop computers, mobile apps, and mobile browsers, respectively.

The results highlight the following facts. First, the number of monthly days with logins increase up to 6 months after the adoption of advice. Economically, the effect is large for the month when advice is implemented, where individuals on average login five more times per month. This increase in attention is still rather large on the following month—with two additional login days. It then decreases to only one additional login day up to the six-month mark. Attention monotonically decreases, with the horizon and the coefficient on for the $35th$ robo-dummy, equals 1.5, suggesting that those who sign up for robo-advice pay attention 1 to 2 days less per month to their investment portfolio. Economically, this is a significant reduction in attention, as the unconditional number of days with logins for non-advised individuals is equal to 4.3. Second, before initiating the process of advice, investors have attention patterns very close to their long-run steady-state as self-directed investors. If anything, they pay slightly less attention than average.

The results in Panels (b) through (d) further allow us to explore the mechanism. The logins from desktop computers and mobile browsers (Subfigure (b) and Subfigure (d)) follow the same decreasing monotonic pattern as the total number of logins, while attention from mobile apps is always higher after signing up for advice. The mobile app is designed in such a way that the user can quickly get an understanding of the status of his/her finances, but it is not a tool when individuals can do extensive research regarding their investment portfolios. The fact that attention through mobile apps stays high after signing up for advice, while attention through the other means decreases and becomes negative over time, suggests two findings. First, after signing up for advice, there is a significant amount of time where individuals monitor and actively try to understand what the advisor is doing—possibly to monitor its performance. Second, as they gain trust individuals stop exerting significant cognitive resources in understanding how the algorithm operates. They instead choose to quickly login to their account using the mobile app to gather information regarding the current value of their portfolio.

The results in Figure 11 use the number of days with logins within a month as a measure of attention. In Figure Online V, we repeat the exercise using the total number of minutes spent on the advisor website as the dependent variable. The results for the minutes of attention are in line with the logins results with two significant differences. First, as shown in Subfigure (a) and (b), the initial increase in attention paid by the investor post signing up for advice is rather short-lived when measured in minutes. On the month of adoption, the investor pays an additional 90 minutes of attention per month, or one and a half hours. However, the difference in attention is already equal to zero two months after adoption. As time goes by, investors tend to spend 30 minutes less on the website, compared to what they did before signing up for advice. Second, the results for attention through the mobile app are different when we work with minutes rather than login. With minutes, we find that attention monotonically decreases and becomes negative after 18 months of adoption. In both cases, the discrepancy between the results that use logins and minutes as measures of attention suggest that, after adopting the service, individuals log in more often, but spend less time on the platform each time they log in.

Overall, the results in this section suggests that adopting advice not only improves individuals investment performance, but it also allows them to decrease the effort they need to exert to manage their investment portfolios. This reduction in attention is not related to an overall reduction in the investors' awareness of their financial condition, because investors tend to login more often to quickly acquire information regarding their portfolio wealth whenever they need to.

6 Adoption and Attrition in Advice

In the previous sections, we showed that those investors who decide to sign up for a hybrid robo-advice service end up benefitting from both an investment performance perspective as well as from attention perspective. Their benchmark-adjusted performance increases and the time investors dedicate to their finances decreases substantially. In this section, we analyze the determinants of adoption and attrition. Our aim is to understand whether the investors who are likely to benefit more from advice are the ones that 1) decide to sign-up for the service, and 2) maintain the service and do not decide to quit the service.

For these results, we exploit the fact that, in our data, we have information regarding all the individuals that were at—some point—interested in signing up for advice, even though only 31% of them end up adopting the service. Our data also contain detailed information regarding investors' attrition, that is, which investors decide to stop using the service.

6.1 Who Signs up for Advice?

In the dataset containing all the investors potentially interested in advice, we have a total of 319,147 individuals with complete information. Unconditionally, we have that 31% of the individuals who consider signing up for advice end up adopting the service. Conditional on initiating the process of enrollment, the probability raises to 39%. Finally, conditional on enrollment in financial advice, the probability of implementing the portfolio allocations is 100%.

In what follows we analyze the determinants of sign-up, from the moment individuals start the process of initiating advice adoption. This is the time where investors are likely to gather precise information about the characteristics of the service. It can therefore help us understand what individual characteristics have an impact on the adoption of advice. This is crucial, because it can help answer the question of whether it is the individuals who need advice the most or the ones who need it the least that sign-up for the service.

We adopt a comprehensive approach where we condition the probability of sign up on a number of market wide and individual characteristics. We measure these quantities on the month in which the investor signs up to initiate the advice process. We estimate a linear probability model of the form:

$$Sign\_Up_i = \alpha + \beta \, \boldsymbol{x}_i + \epsilon_i,$$

where $Sign\_Up_i$ is an indicator variable equals to 0 if the investor does not adopt advice and 1 if it does. The vector $\boldsymbol{x}_i$ contains the following market-wide regressors: $Market\_Ret$ and $Market\_Var$ are market-wide returns and volatility, respectively. It also includes the following individual specific returns and portfolio-characteristic variables: $Investor\_Ret$ is the investor's monthly return; $Realized\_Var$ is the monthly realized variance; $Percent\_ETF$ is the client percentage of wealth in ETF; $Percent\_MF$ is the percentage of wealth in mutual funds; $Percent\_Stock$ is the percentage of wealth in individual stocks; $Expense\_Ratio$ is the expense ratio paid on the mutual funds held;

Percent_International is the percentage of wealth in international stocks or bonds; *Client_Wealth* is overall client wealth of the investor; *Equity_Share* and *Cash_Share* are the investors' portfolio shares in equities and cash (including money market mutual funds), respectively. The last two regressors included are *Human_Advisor_Scheduled_Appointments* and *Investor_Scheduled_Appointments*. These appointments include appointments categorized as client-facing, individual consults, plan delivery, ad-hoc meetings, check-in, and plan delivery. All non-binary regressors are standardized so that they have unit variance. This helps with the interpretation of the results.

The results are reported in the first column of Table 7. The coefficients on the market return and volatility variables suggest aggregate market returns do not affect investors' decision to sign-up for advice, but market volatility does. The higher the market uncertainty investors face when they initiate the process of advice, the more likely they are to sign-up.

With respect to the question of whether it is the individuals who benefit the most from advice to be the ones who sign up for it, we find the answer is yes, with a couple of exceptions. The higher the individual investor portfolio volatility, the higher the probability of eventually signing up for the advice service. Individuals are likely to sign-up for advice also when their portfolio is doing poorly, as evidenced by the negative and significant coefficient on *Investor_Ret*.

Also, it is those individuals holding mutual funds with high expense ratios, relatively high stock holdings and those who have little international diversification the ones that sign-up for robo-advising. Economically, the magnitudes for some of these regressors are rather large. For example, a standard deviation increase in investors' expense ratios increases investor's probability to sign up for advice by 5 percentage points. Given a constant of 21%, the increase implies approximately a 5/21=25% increase in the probability of signing up for advice. On the other hand, mutual fund holdings or cash holdings do not play a large role in the decision to sign up for the service.

The last individual characteristics we consider are investor wealth and exposure to equities. Both are negatively related to the probability of eventually adopting robo-advising. The coefficient on *log_client_wealth* is particularly large. In the cross-section, a standard deviation increase in wealth lowers the probability of signing up for advice by 3.6 percentage points. Once again, economically the coefficients are very large. Given a constant of 21%, the increase implies approximately a 3.6/21=17% increase in the probability of signing up for advice.

The last variables we consider are clients' interactions with advisors. This is associated more with the advisors' characteristics rather than clients' characteristics. Investors' scheduled appointments do not seem to have an impact on the probability to sign-up. What seems to have a major impact, on the other hand, are advisors' scheduled appointments. Having an appointment scheduled on the very same month of initiating enrollment into financial advice increases the probability of adopting financial advice by seven percentage points, indicating a very large effect of the human touch on the adoption of advice.

In the second column of Table 7, we repeat the analysis using a Logit specification instead of a linear probability model. The results are virtually identical, both from an economic and statistical perspective.

Overall, the results suggest that it is the individuals that benefit the most from robo-advising—that is— those who have low international diversification, have high expense ratios, and have high portfolio volatility. The results also suggest that immediacy in the scheduling of an appointment with a human advisor plays a very important role in the probability of investor signing up for advice.

6.2 Attrition in Financial Advice

In this section, we study the other side of the decision to sign up for advice, i.e., the choice of quitting the advice service.

To motivate the analysis, we start by presenting simple Kaplan Meier estimates of investor attrition, starting the computations from the day of enrollment. In Figure 12, we present non-parametric estimates of the percentage of investors maintaining their subscription to financial advice as a function of the years they have been subscribed to the service. Subfigure (a) compares attrition among male and female investors. The blue line reports results for male investors and the red line the ones for female investors. In both cases, we report mean estimates and 95% confidence intervals. The plot reveals two facts. First, male and female investors behave similarly, when it comes to quitting advice. Second, the attrition rate is less than 5% per year. In subfigure (b), we divide the investors in long-tenure (blue line and confidence interval) and short-tenure (red line and confidence interval) investors. Long-tenure investors are less likely to quit advice by 5 percentage points at the four-year mark, suggesting the importance of trust and familiarity in not quitting the service.

In subfigure (c) the blue line reports results for "quick to enroll" investors and the red line results for "slow to enroll" investors. This is simply measured as the time elapsed between initiating the on-boarding process into financial advice and the enrollment into the service. Interestingly, we do not find differences in attrition between those who quickly complete the enrollment procedure and those that don't.

Finally, in subfigure (d) the blue line reports results for Level 1 investors, who have \$500,000 at most in assets under management (AUM); the red line reports results for Level 2 investors, who have between \$500,000 and \$1 million in AUM (red line); and the green line reports results for Level 3 investors, who have between \$1 million and \$5 millions in AUM. The attrition across the three groups is markedly different. Almost 90% of the Level 3 investors remain invested after four years. The percentage equals 83% and 76% for Level 2 and Level 1 investors, respectively. The differences in behavior could be due to a number of factors. First, higher wealth individuals are given a higher level of service. For example, they have a dedicated advisor rather than being serviced by a pool of advisors. Second, the service is cheaper for wealthier individuals. The fees amount to 30 basis points for investable assets below \$500,000, they are 25 basis points for investable assets between \$500,000 and \$1 million and 20 basis points for investable assets between \$1 million and \$5 millions.

Next, we perform a more comprehensive analysis that parallels the adoption results, but focuses on attrition. We estimate a Cox Proportional Hazard model of the form:

$$\lambda(t|\boldsymbol{x}_i) = \lambda_0(t) \exp(\boldsymbol{x}_i \times \boldsymbol{\beta}),$$

where $\lambda(t|\boldsymbol{x}_i)$ is the hazard function at time t for investor i with covariate vector (explanatory variables) $\boldsymbol{x}_i$. The vector $\boldsymbol{x}_i$ contains the same regressors described in Section 6.1.

The hazard ratios are reported in Table 8. The first column reports the baseline results. The second column includes the two client segmentation dummies "Level 1" and "Level 2" investors.

We highlight several findings. First, very few covariates measured at the time of sign-up explain attrition. Among the exceptions we find overall market variance ($Market\_Var$). Those investors who sign up in periods of high market volatility are approximately 10% more reluctant to quit the advice service—hazard ratio equal to 0.907, significant at the 1% level. We find a similar effect for expense ratios ($Expense\_Ratio$). Those paying high expense ratios before signing up for advice are 10% more reluctant to quit advice after they enroll—hazard ratio equal to 0.910 significant at the 1% level. We find an opposite effect for equity share that is instead positively related to attrition. Those with a higher equity share at the time of sign-up are 10% more likely to quit advice—hazard ratio equal to 1.104, significant at the 1% level.

Advisor-scheduled appointments are not important determinants of attrition, while investor-scheduled appointments are. The more investors reach out to human advisors, the less likely they are to quit the service. Finally, we find that the Level 1 and Level 2 service dummies are extremely significant. The base case are Level 3 investors, which are the least reluctant to quit advice, as shown in Figure 12. The results in 8 show that Level 1 investors are 100% more likely to quit advice, compared to Level 3 investors. Level 2 investors are 40% more likely to quit advice, compared to Level 3 investors.

The results for client wealth (regressor $Client\_Wealth$) highlight the importance of controlling for the type of service investors sign-up. In the first column, the coefficient on $Client\_Wealth$ is insignificantly different from zero. Once we control for the type service investors receive, the coefficient on $Client\_Wealth$ becomes positive and economically significant. An increase in income increases the probability of quitting advice by 22%—hazard ratio equal to 1.220 significant at the 1% level.

Overall, the negative relation between expense ratios and probability of quitting advice supports the notion that those investors who benefit the most from advice are the ones who are less likely to quit the service.

7 Conclusions

We study the effects of a large U.S. hybrid robo-advisor on the portfolios of previously self-directed investors.

Across all investors, robo-advising reduces investors' holdings in money market mutual funds and increases bond holdings. It also reduces idiosyncratic risk by lowering the holdings of individual stocks and US active mutual funds and raising exposure to low-cost indexed mutual funds. It further eliminates home bias by significantly increasing international equity and fixed income diversification. Over our relatively short sample period, it increases investors' overall risk-adjusted performance.

We use a machine learning algorithm, known as Boosted Regression Trees (BRT), to explain the cross-sectional variation in the effects of advice on portfolio allocations and performance. Investors who benefit from advice are those with little investment experience, those with prior high cash holdings, and those with high trading volume before adopting advice. Investors with little mutual fund holdings and clients invested in high-fee active mutual funds also display significant performance gains.

In the final part of the paper, we study the determinants of investors' sign-up and attrition. The evidence suggests that those investors who benefit more from advice are also more likely to sign up and less likely to quit the service.

References

Abel, A. B., J. C. Eberly, and S. Panageas. 2007. Optimal inattention to the stock market. *American Economic Review* 92:244–49.

———. 2013. Optimal inattention to the stock market with information costs and transactions costs. *Econometrica* 81:1455–81.

Aghion, P., and P. Howitt. 1992. A model of growth through creative destruction. *Econometrica* 60:323–51.

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

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

Bhattacharya, U., A. Hackethal, S. Kaesler, B. Loos, and S. Meyer. 2012. Is unbiased financial advice to retail investors sufficient? answers from a large field study. *The Review of Financial Studies* 25:975–1032.

Breiman, L. 1996. Bagging predictors. *Machine learning* 24:123–40.

Breiman, L., J. H. Friedman, R. A. Olshen, and C. J. Stone. 1984. *Classification and regression trees.* Statistics/Probability Series. Belmont, California, U.S.A.: Wadsworth Publishing Company.

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

Comin, D., and M. Mestieri. 2014. Technology diffusion: Measurement, causes, and consequences. In P. Aghion and S. N. Durlauf, eds., *Handbook of Economic Growth*, vol. 2 of *Handbook of Economic Growth*, 565 – 622. Elsevier.

D'Acunto, F. 2018a. Identity, beliefs, and choice under risk. *Working Paper* .

———. 2018b. Tear down this wall street: Anti-market rhetoric, motivated beliefs, and investment. *Working Paper* .

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

D'Acunto, F., M. Prokopczuk, and M. Weber. Forthcoming. Historical antisemitism, ethnic specialization, and financial development. *Review of Economic Studies* .

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

Friedman, J., T. Hastie, and R. Tibshirani. 2001. *The elements of statistical learning*, vol. 1. Springer series in statistics New York, NY, USA:.

Friedman, J. H. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* 1189–232.

———. 2002. Stochastic gradient boosting. *Computational Statistics & Data Analysis* 38:367–78.

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.

Guiso, L., M. Haliassos, and T. Jappelli. 2003. Household stockholding in europe: where do we stand and where do we go? *Economic Policy* 18:123–70.

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

Kim, H. H., R. Maurer, and O. S. Mitchell. 2016. Time is money: Rational life cycle inertia and the delegation of investment management. *Journal of financial economics* 121:427–47.

Linnainmaa, J. T., B. Melzer, and A. Previtero. forthcoming. The misguided beliefs of financial advisors. *Journal of Finance* .

———. Forthcoming. The misguided beliefs of financial advisors. *The Journal of Finance* .

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

Rapach, D. E., J. K. Strauss, and G. Zhou. 2010. Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *The Review of Financial Studies* 23:821–62.

Reher, M., and S. Sokolinski. 2020. Does fintech democratize investing? *Unpublished Working Paper* .

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

Romer, P. M. 1990. Endogenous technological change. *Journal of political Economy* 98:S71–S102.

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

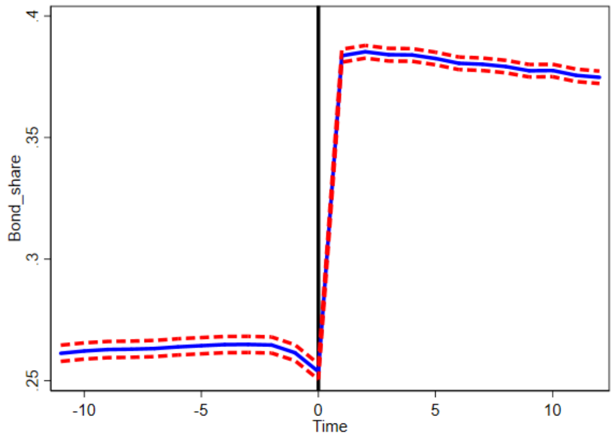

(a) Percentage of Bonds

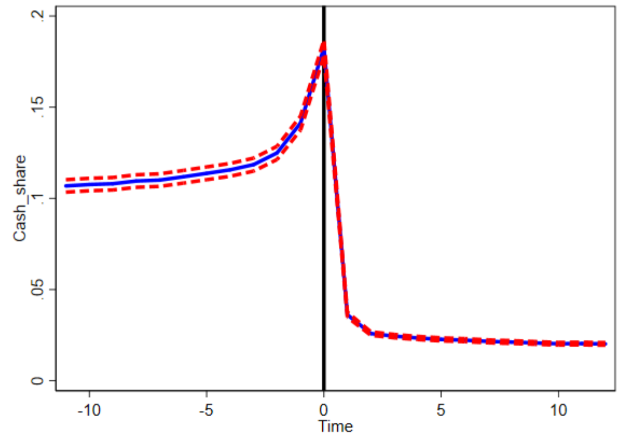

(b) Percentage of Cash

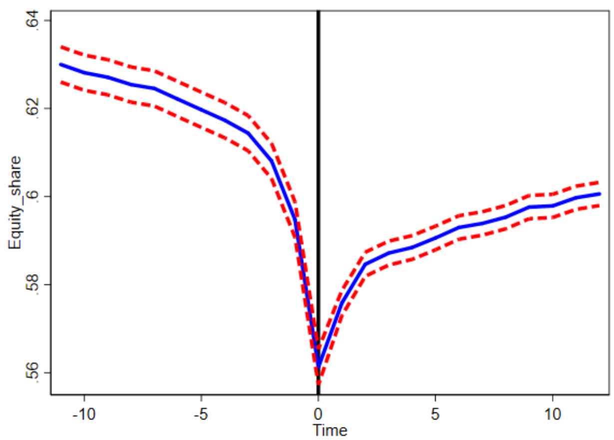

(c) Percentage of Equities

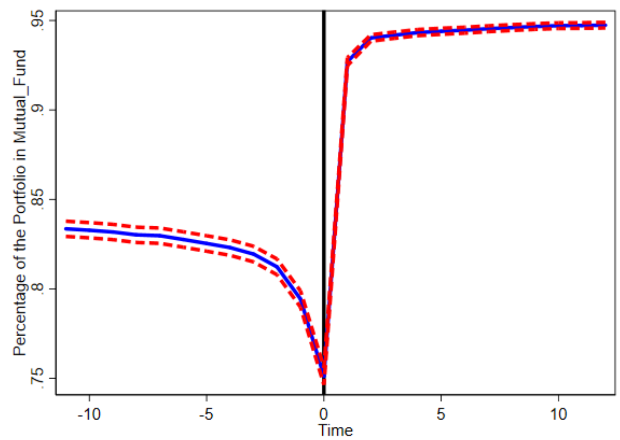

(d) Percentage in Mutual Funds

Figure 1. Portfolio Allocation Before and After Advice: All Investors. This figure reports results for investor portfolio characteristics before and after signing up for robo-advice. The results are computed at the investor level and include all account types, that is, taxable and non-taxable (IRA) accounts. Subfigures (a), (b), (c), and (d) report results for the percentage of wealth held—directly or through mutual funds—in bonds, cash (including money market mutual funds), equities and mutual funds. In each subfigure, time "0" represent the month before investors sign up for advice. Results are computed using only investors that are in the sample for at least twelve months before and after signing up for advice. The blue line denotes average values, while the red dashed lines are 95% confidence intervals.

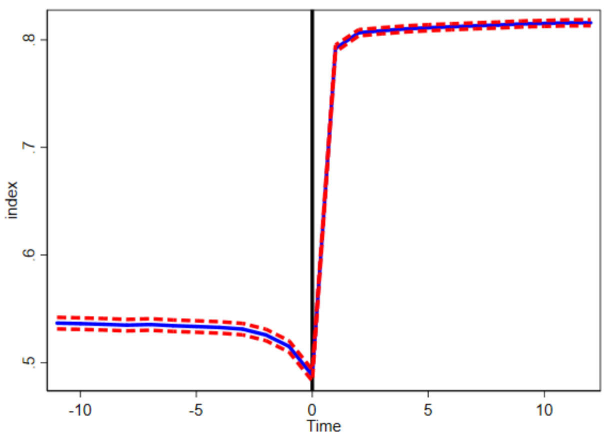

(a) Percentage of Indexed Funds

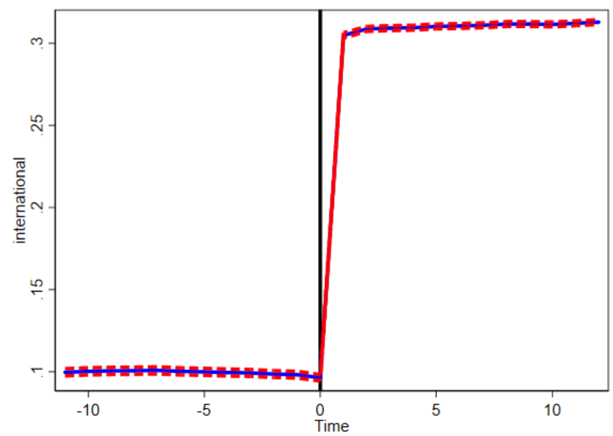

(b) Percentage of International Funds

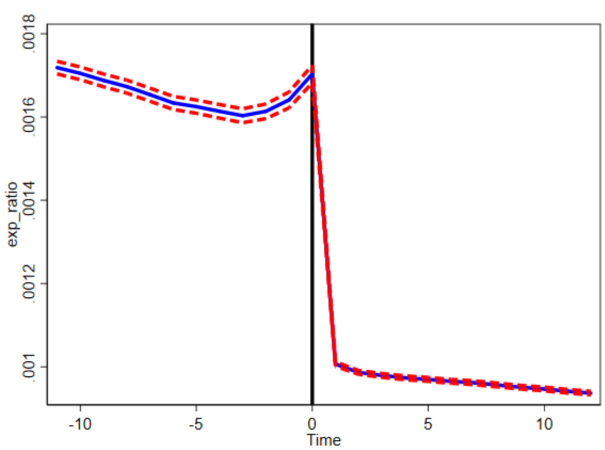

(c) Expense Ratios

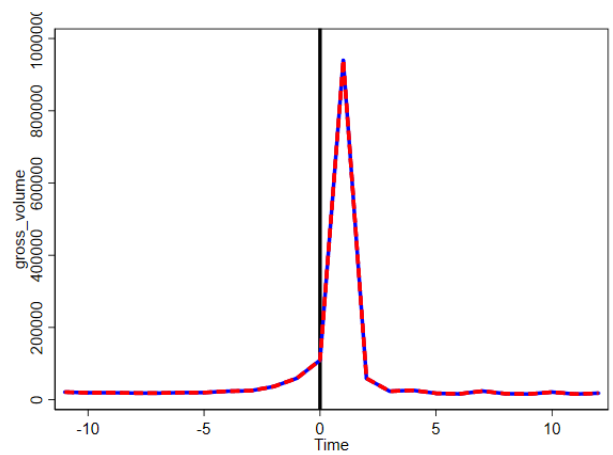

(d) Trading Volume

Figure 2. Indexation, International Diversification and Fees Before and After Advice: All Investors. This figure reports results for investor portfolio characteristics before and after signing up for robo-advice. The results are computed at the investor level and include all account types, that is, taxable and non-taxable (IRA) accounts. Subfigure (a) reports results for the percentage of mutual fund wealth invested in indexed funds; Subfigure (b) the percentage of mutual fund wealth invested in global or international funds—identified using the Lipper mutual fund classification; Subfigure (c), the value-weighted expense ratio charged by the mutual funds held by the account-holders. Finally, Subfigure (d) shows results for the monthly trading volume, in US dollars. In each subfigure, time "0" represent the month before investors sign up for advice. Results are computed using only investors that are in the sample for at least twelve months before and after signing up for advice. The blue line denotes average values, while the red dashed lines are 95% confidence intervals.

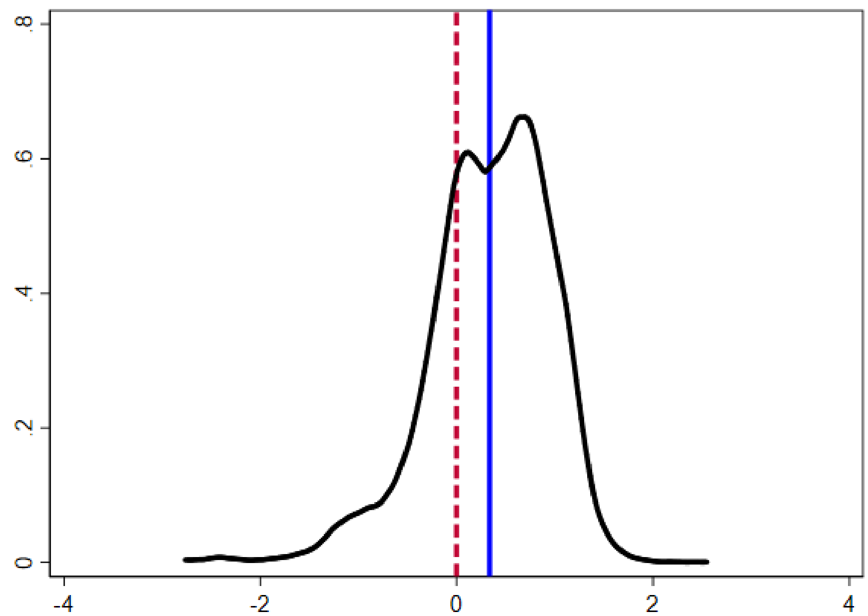

Figure 3. Pre- and post-Advice Performance Difference Distribution: Investor Level. This figure reports the change in investment performance across Vanguard investors after signing up for the Personal Advisor Service (PAS) from Vanguard, compared to before signing up for PAS. The results are computed at the investor level and are computed across all accounts, As a measure of performance, we use the abnormal Sharpe Ratio, that is, the difference between the realized Sharpe ratio of each account-holder and the realized Sharpe ratio of the market portfolio, where the latter is computed as the value-weighted returns on the NYSE/AMEX/NASDAQ CRSP portfolio. For each account-holder, portfolio returns and volatilities are computed using beginning-of-month investment holdings. Furthermore, portfolio volatilities are computed as realized volatilities using squared daily returns. The figure uses all the accounts that have been in the sample for at least 6 months before and after signing up for advice and computes abnormal Sharpe ratios using 9-month windows. The black solid line denotes the density of the difference in post- and pre-advice performance, computed at the investor level. The red dashed line marks the value "0" of the x-axis, while the blue solid line denotes the mean of the distribution.

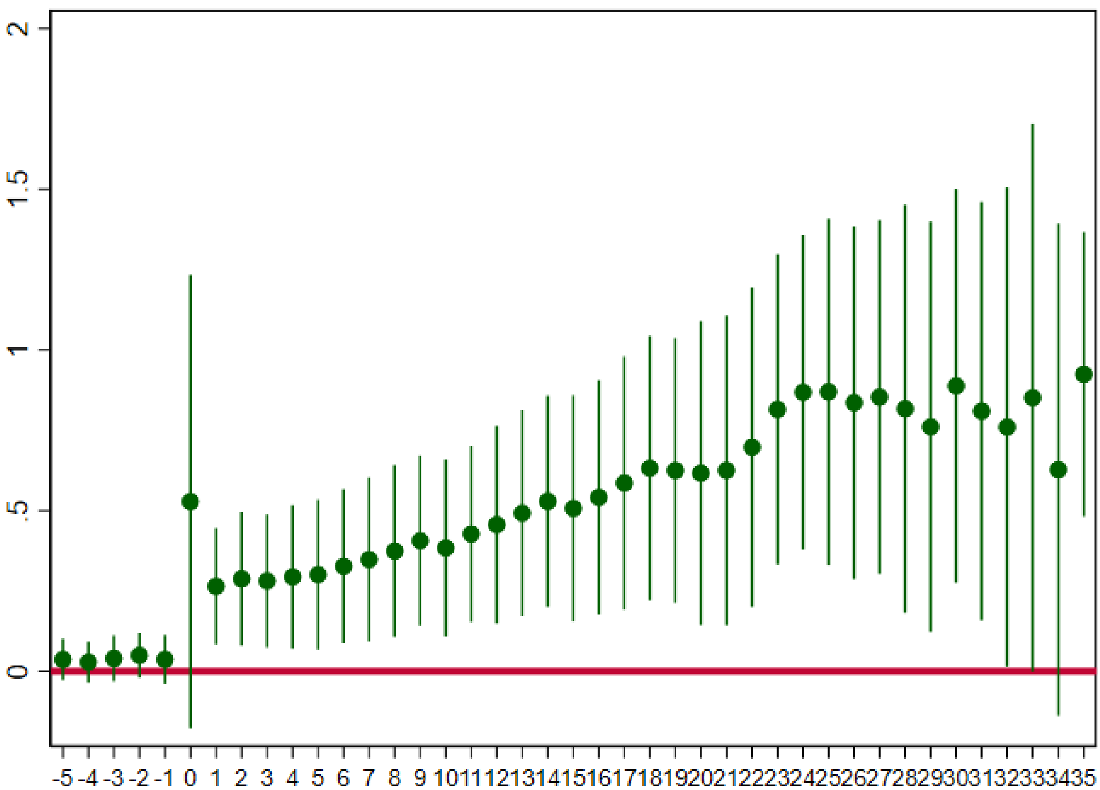

Figure 4. Panel Regressions Relating Robo-Advising Adoption and Performance. This figure reports results relating robo-advising adoption and investment performance. The regression estimated is:

$$Sharpe_{i,t} = \alpha_i + \beta_t + \sum_{j=-5}^{35} \gamma_j \; ROBO_{i,j,t} + \epsilon_{i,t},$$

where $Sharpe_{i,t}$ represents the monthly Sharpe ratio, α_i denotes individual fixed effects, β_t monthly time-effects, and $ROBO_{i,j,t}$ is is a dummy variable equal to zero for every month, except for the j-th month before and after adoption. The $0 - th$ month is the one when the robo-advisor is adopted, negative values of j refer to the months before advice is adopted, and positive values of j refer to months after robo-advice is adopted. The figure plots the γ_j coefficients and their 95% confidence intervals, computed using double-clustered standard errors.

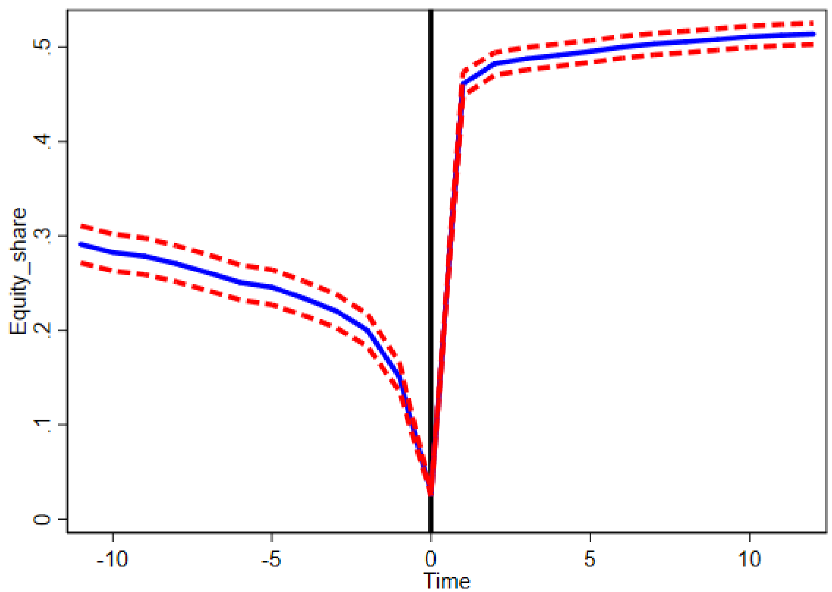

(a) Low Equity Share at PAS Sign-up

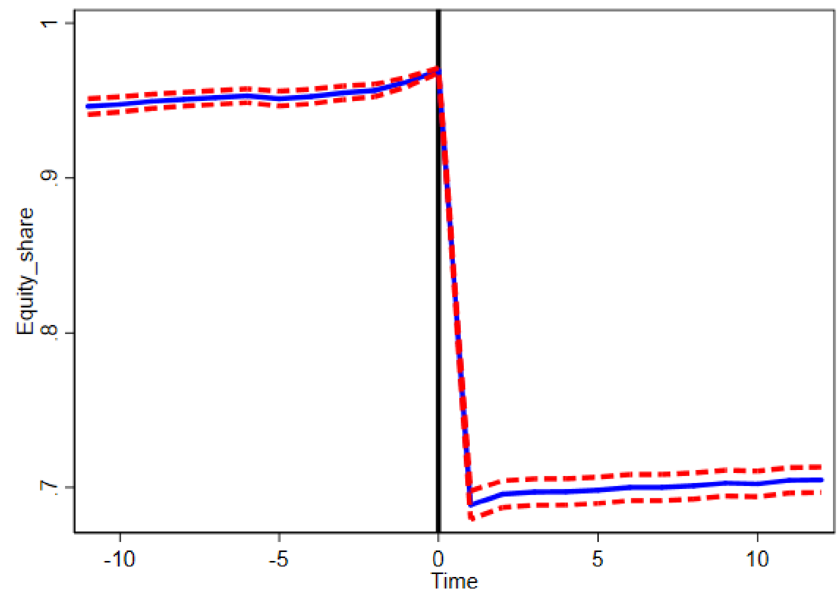

(b) High Equity Share at PAS Sign-up

Figure 5. Portfolio Changes for Investors with Low and High Equity Shares Pre-advice. This figure reports results for the changes in equity shares for Vanguard investors, before and after signing up for the Personal Advisor Service (PAS) from Vanguard. The results are computed at the investor level and include all account types, that is, taxable and non-taxable (IRA) accounts. Subfigure (a) reports results for investors with less than 10% of their portfolio in equities pre-advice Subfigure (b) reports results for investors with more than 90% of their portfolio in equities post-advice. In each subfigure, time "0" represent the month before investors sign up for advice. Results are computed using only investors that are in the sample for at least twelve months before and after signing up for advice. The blue line denotes average values, while the red dashed lines are 95% confidence intervals.

Panel A. Univariate Partial Dependence Plots

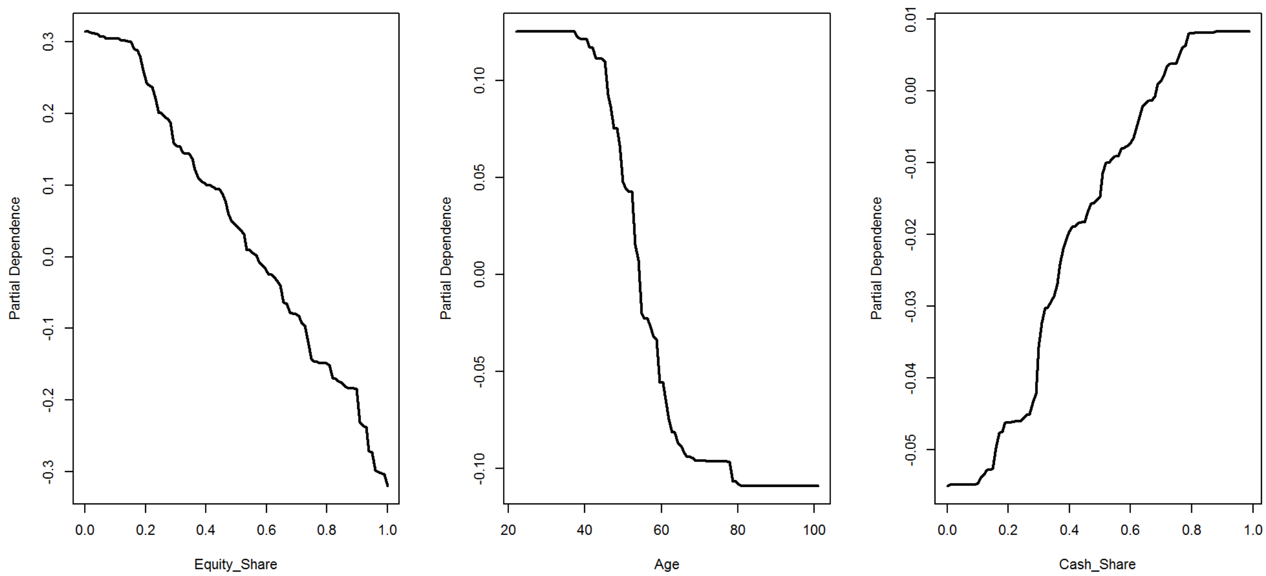

Panel B. Bivariate Partial Dependence Plots

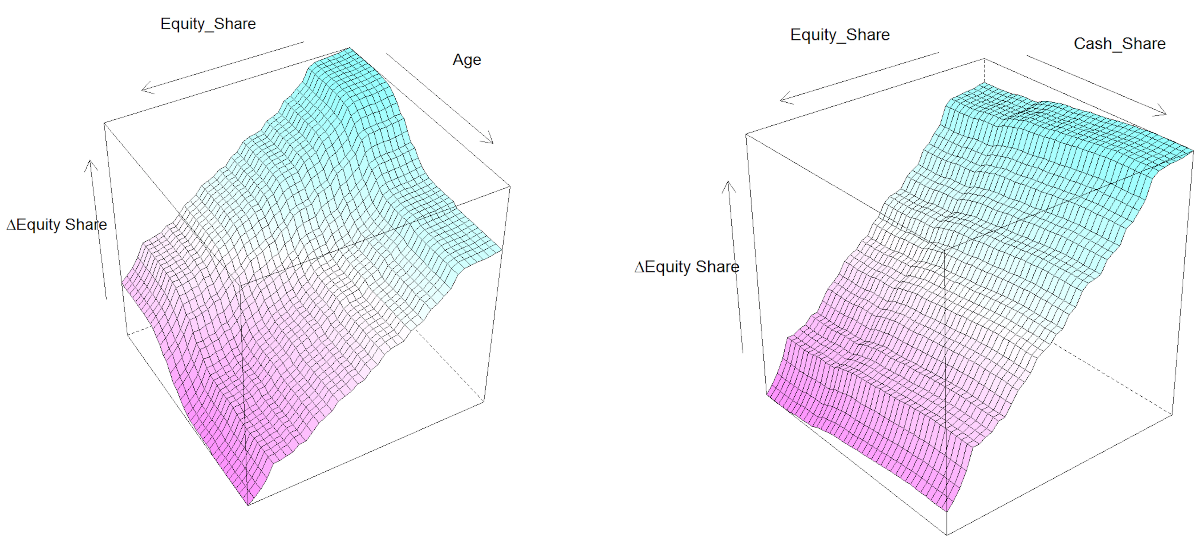

Figure 6. Partial Dependence Plots for the Change in Equity Share Post-advice. This figure presents the partial dependence plots for the change in equity share as a function of a total of 15 regressors described in Section 4.1. In Panel A of the figure, we report partial dependence plots for the three predictor variables with the highest relative influence: $PctEquityShare$, the share of wealth in Bonds (relative influence of 81.9%); Age, the age of the investor (relative influence of 15.6%); and $PctCashShare$, the share of wealth in cash (relative influence of 2.1%). In Panel B, we report bivariate partial dependence plots for $PctEquityShare$ and Age, and for $PctEquityShare$ and $PctCashShare$. The horizontal axis covers the sample support of each predictor variable, while the vertical axis tracks the change in the equity share as a function of each individual predictor variable.

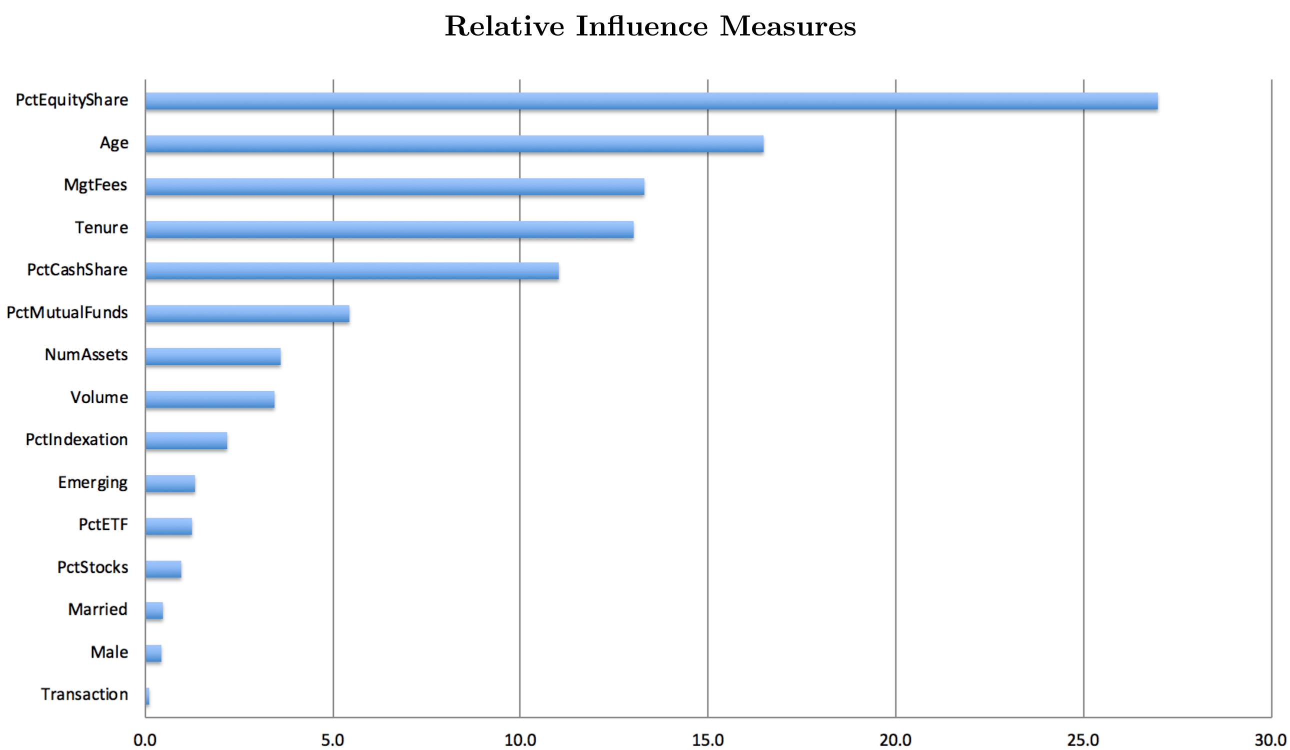

Figure 7. Relative Influence Plots for the Change in Performance Post-advice. This figure presents the relative influence plots for the change performance post adopting advice as a function of a total of 15 regressors described in Section 4.1. The relative influence value associated with each regressor corresponds the relative importance of the covariate in explaining the changes in performance post-adoption of advice. By construction, the sum of the relative influences across all the covariates sums to 100.

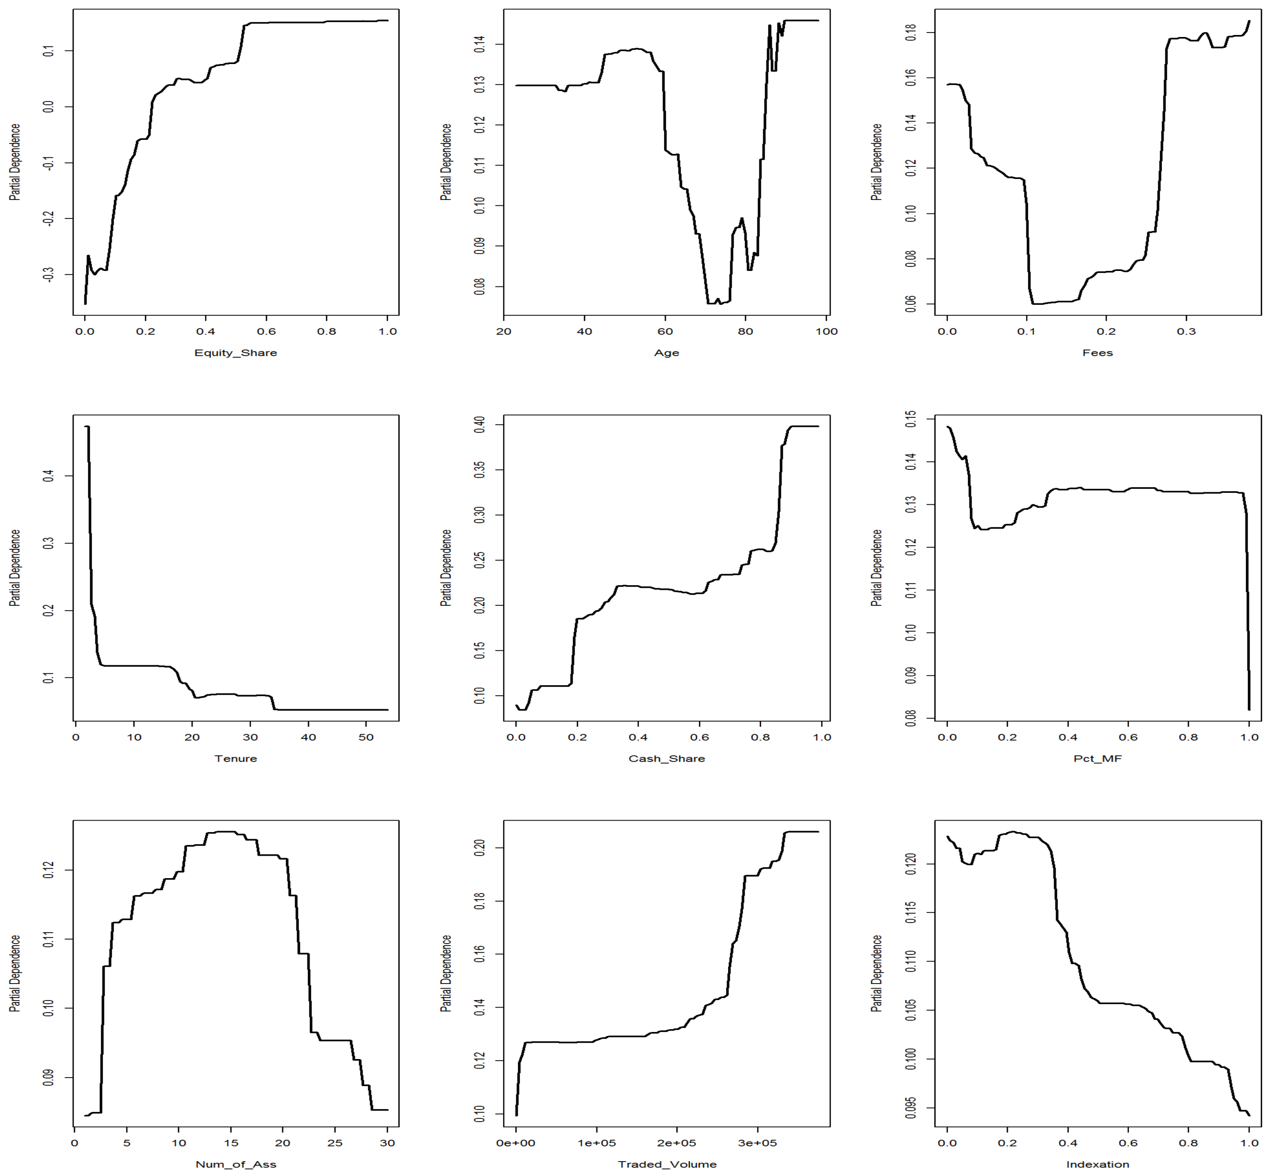

Figure 8. Partial Dependence Plots for Investment Performance Changes Post-advice. This figure presents the partial dependence plots for the change in abnormal Sharpe Ratio as a function of a total of 15 regressors described in Section 4.1. We report partial dependence plots for the 9 predictor variables with the highest relative influence: $PctEquityShare$, the share of wealth in Equities (relative influence of 27%); Age, the age of the investor (relative influence of 16.5%); $MgtFees$, the value-weighted management fees charged by the mutual funds held by the investor (relative influence of 13.3%); $Tenure$, the tenure of the investor as of December 2017 (relative influence of 13.0%); $PctCashShare$, the percentage of wealth money market mutual funds—directly or through mutual funds (relative influence of 11.0%); $PctMutualFunds$, the percentage of wealth directly invested in mutual funds (relative influence of 5.4%); $NumAssets$, the number of assets held by the investor across accounts (relative influence of 3.6%); $Volume$, the volume (in US dollars) traded by the investors over the month before signing up for advice (relative influence of 3.4%); $PctIndex$, the percentage of mutual fund wealth invested in index funds (relative influence of 2.2%). The horizontal axis covers the sample support of each predictor variable, while the vertical axis tracks the change in the equity share as a function of each individual predictor variable.

Performance Across Boosting Iterations

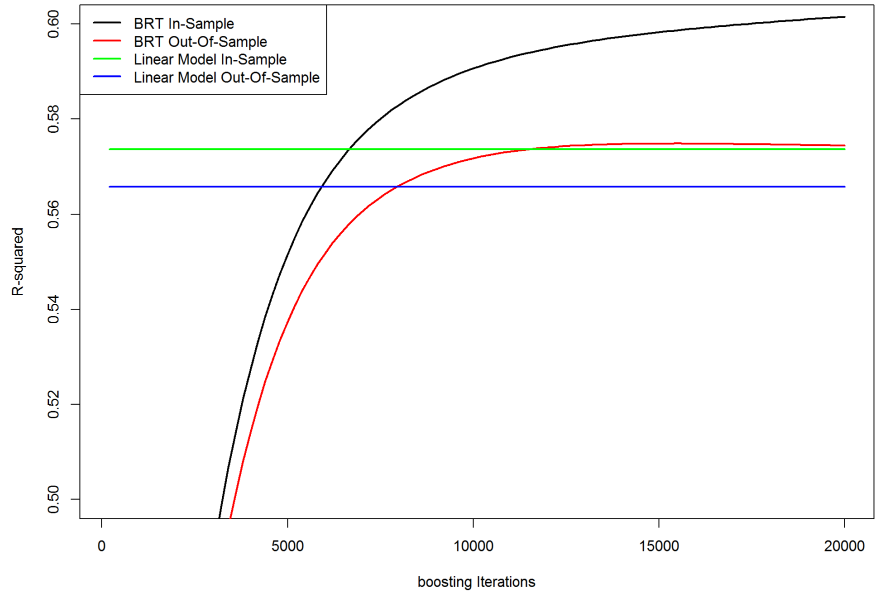

Performance across Monte Carlo Samples

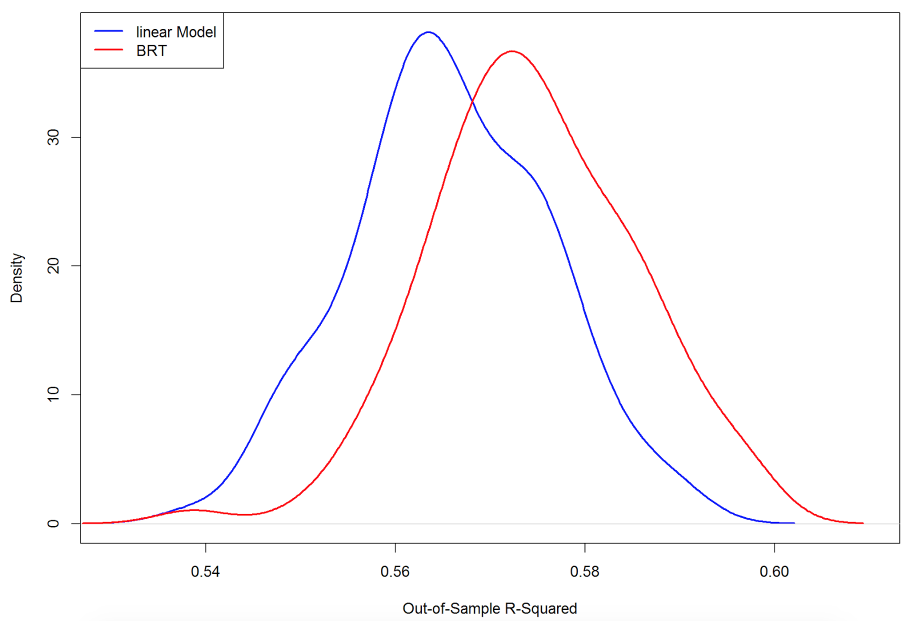

Figure 9. In- and Out-of-Sample Average BRT Performance Across Boosting Iterations and Monte Carlo Samples for Portfolio Changes Before and After Signing up for Advice. This figure plots in, Subfigure A, the in- and out-of-sample performance, across boosting iterations, for a Boosted Regression Trees model and a linear regression model that uses the same covariates. The BRT in-sample performance is denoted by a black line; the BRT out-of-sample performance is denoted by a red line; the linear model in-sample performance is denoted by a green line; and the linear model out-of-sample performance is denoted by a blue line. Subfigure B plots densities of out-of-sample performance for a Boosted Regression Trees model with 20,000 boosting iterations and a linear regression model that uses the same covariates. The BRT performance is denoted by a red line while the linear model performance is denoted by a blue line.

Performance Across Boosting Iterations

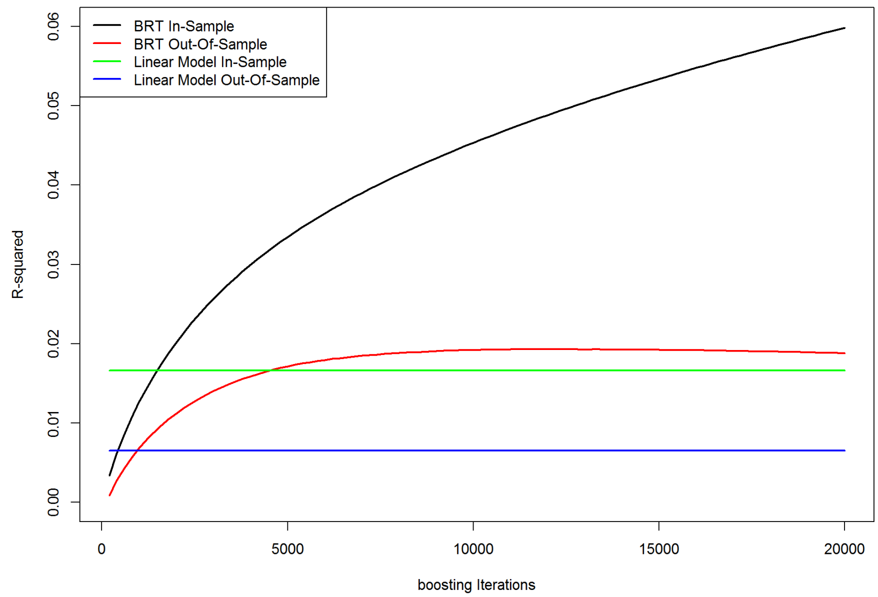

Performance across Monte Carlo Samples

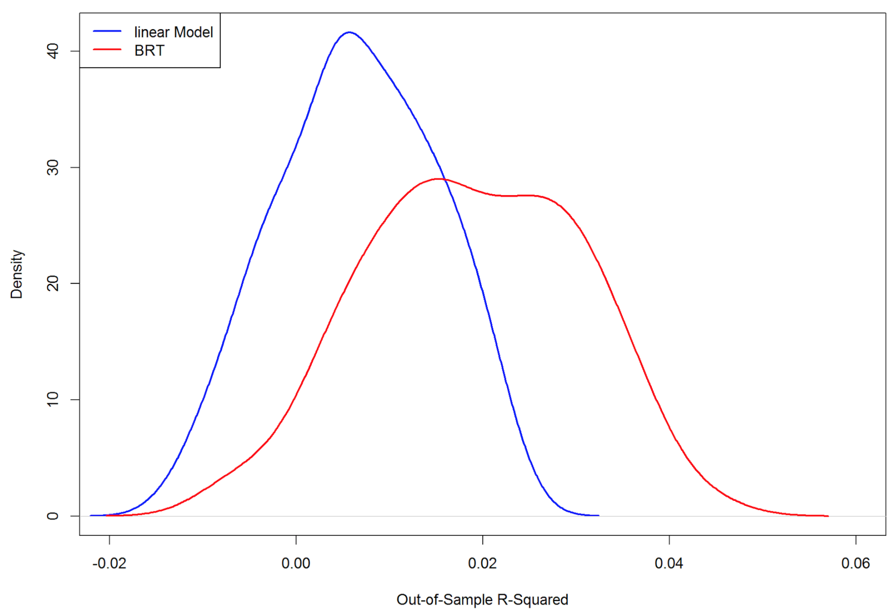

Figure 10. In- and Out-of-Sample Average BRT Performance Across Boosting Iterations and Monte Carlo Samples for Investment Performance Changes Before and After Signing up for Advice. This figure plots in, Subfigure A, the in- and out-of-sample performance, across boosting iterations, for a Boosted Regression Trees model and a linear regression model that uses the same covariates. The BRT in-sample performance is denoted by a black line; the BRT out-of-sample performance is denoted by a red line; the linear model in-sample performance is denoted by a green line; and the linear model out-of-sample performance is denoted by a blue line. Subfigure B plots densities of out-of-sample performance for a Boosted Regression Trees model with 20,000 boosting iterations and a linear regression model that uses the same covariates. The BRT performance is denoted by a red line while the linear model performance is denoted by a blue line.

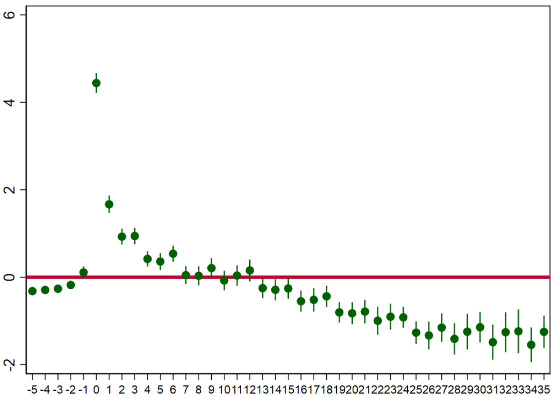

(a) Total

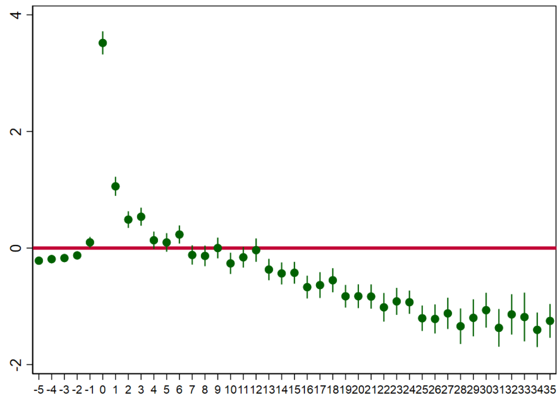

(b) Attention Through Desktop Computer

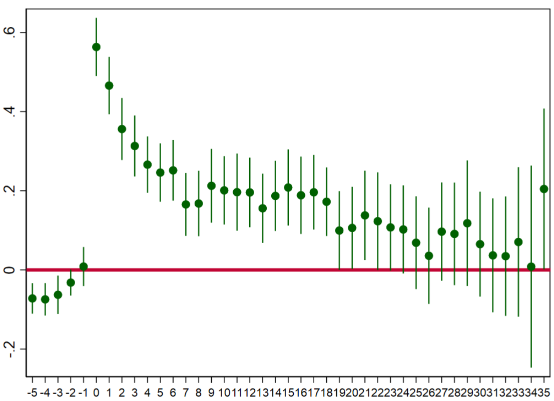

(c) Attention Through Mobile App

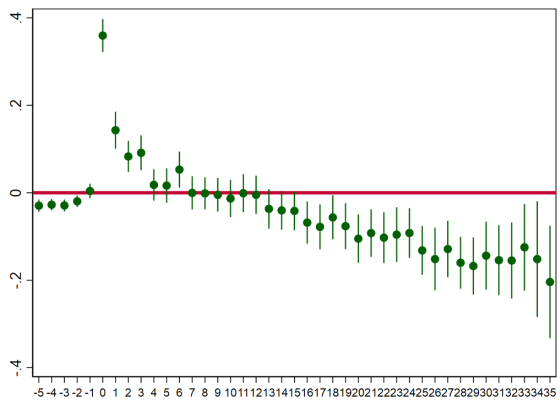

(d) Attention Through Mobile Browser

Figure 11. Panel Regressions Relating Robo-Advising Adoption and Investors' Login-Attention. This figure reports results relating robo-advising adoption and investment performance. The regression estimated is:

$$Attention_{i,t} = \alpha_i + \beta_t + \sum_{j=-5}^{35} \gamma_j \ ROBO_{i,j,t} + \delta \ X_{i,t} + \epsilon_{i,t},$$

where $Attention_{i,t}$ denotes the number of days with logins by investor i on month t, α_i denotes individual fixed effects, β_t monthly time-effects, and $ROBO_{i,j,t}$ is a dummy variable equal to zero for every month, except for the j-th month before robo-advising enrollment is initiated and after robo-advising is implemented. The $0-th$ month is the one when the robo-advisor is implemented, negative values of j refer to the months before the process of robo-advising is initiated. We stop at the $35-th$ month post-adoption as it is the longest horizon for which we have robo-advised clients. Note that in these results we do not include dummies for the periods of robo-advising enrollment as the enrollment process range between one to several months across individuals. The regressor $X_{i,t}$ contains investors' portfolio return and risk, which do not have an economically significant effect on attention. The figure is composed of four subfigures. Subfigure (a) reports the results for the total number of logins; subfigures (b) through (d) report results for the logins through desktop computers, mobile apps and mobile browsers, respectively. Each subfigure plots the γ_j coefficients and their 95% confidence intervals, computed using double-clustered standard errors.

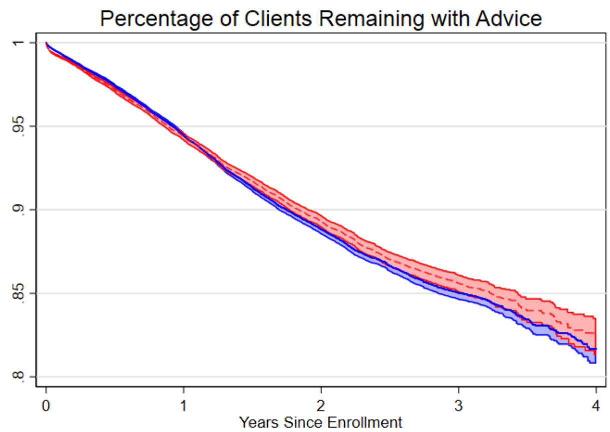

(a) Red: Female Investors; Blue: Male Investors

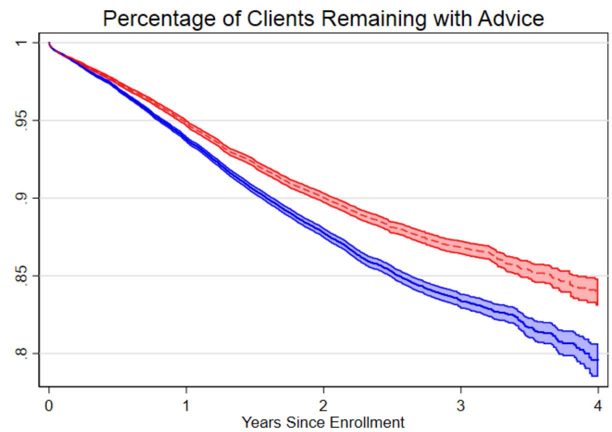

(b) Red: Long-tenure Investors; Blue: Short-tenure Investors

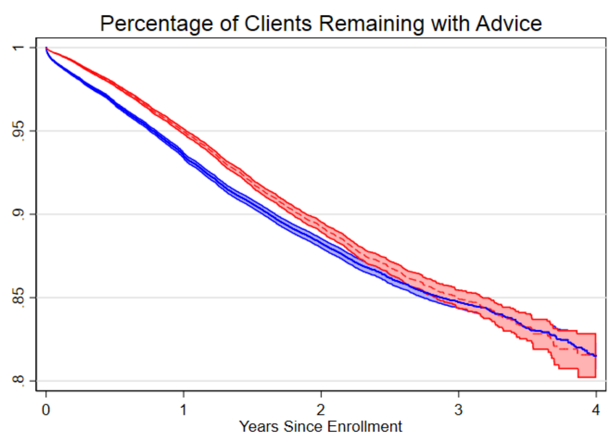

(c) Red: Slow to Enroll; Blue: Quick to Enroll

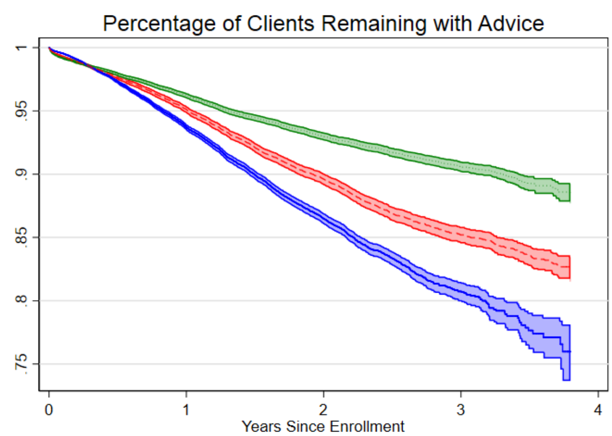

(d) Green: Level 3; Red: Level 2; Blue: Level 1

Figure 12. Kaplan Meier Estimates of the Investor Attrition in Advice. These plots present investor attrition from enrollment into financial advice. Each line within each plot presents estimates of the percentage of clients remaining subscribed into financial advice as a function of the years after enrollment—together with its 95% confidence interval. In subfigure (a), the blue line reports results for male investors and the red line the ones for female investors. In subfigure (b), the blue line reports results for long-tenure investors and the red line the ones for short-tenure investors. In subfigure (c) the blue line reports results for "quick to enroll" investors and the red line results for slow to enroll investors. Finally, in subfigure (d) the blue line reports results for Level 1 investors, who have \$500K at most in AUM; the red line reports results for Level 2 investors, who have between \$500K and \$1M in AUM (red line); and the green line reports results for Level 3, who have between \$1M and \$5M dollars in AUM.

Table 1. Demographic and Portfolio Characteristics of Robo-advised investors before Sign-up.

Panel A. Demographic Characteristics

	N	mean	sd	p10	p25	p50	p75	p90
Age	80,690	63.22	12.80	45.00	56.00	65.00	71.00	78.00
Male	82,526	0.53	0.50	0.00	0.00	1.00	1.00	1.00
Married	82,526	0.69	0.46	0.00	0.00	1.00	1.00	1.00
Tenure	82,498	14.18	9.30	2.00	5.08	14.17	20.67	26.08

Panel B. Portfolio Allocation

	N	mean	sd	p10	p25	p50	p75	p90
Wealth ($)	82,526	588,246	832,297	42,604	107,320	282,450	690,604	1,440,446
NumAssets	82,526	7.79	7.95	1.00	2.00	5.00	10.00	17.00
PctEquityShare	81,869	0.54	0.31	0.00	0.33	0.59	0.78	0.95
PctBondShare	81,869	0.24	0.23	0.00	0.00	0.20	0.40	0.58
PctCashShare	81,869	0.22	0.34	0.00	0.00	0.02	0.28	1.00
PctMutualFunds	82,364	0.72	0.37	0.00	0.48	0.94	1.00	1.00
PctCash	82,364	0.20	0.34	0.00	0.00	0.01	0.23	0.98
PctStocks	82,364	0.03	0.10	0.00	0.00	0.00	0.00	0.11
PctETF	82,364	0.03	0.10	0.00	0.00	0.00	0.00	0.05
PctIndex	82,523	0.47	0.37	0.00	0.05	0.46	0.82	1.00
PctInternational	77,083	0.10	0.14	0.00	0.00	0.02	0.17	0.29
PctEmerging	77,083	0.00	0.02	0.00	0.00	0.00	0.00	0.00

Panel C. Transactions and Fees

	N	mean	sd	p10	p25	p50	p75	p90
MgtFee	76,986	0.14	0.12	0.03	0.06	0.11	0.17	0.28
ExpRatio	73,582	0.19	0.17	0.07	0.09	0.14	0.21	0.36
TurnRatio	72,930	0.32	0.26	0.07	0.14	0.25	0.40	0.63
Transaction	82,526	3.31	6.55	0.00	0.00	1.00	3.00	10.00
Volume ($)	82,526	85,247	226,359	0.00	0.00	227	30,000	258,298

This table reports demographic characteristics and portfolio allocation behavior of investors the month before 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: *Age*, the age of the investor as of December 2017; *Male*, whether the investor is male; *Married*, whether he/she is married; *Tenure*, the tenure of the investor as of December 2017. Panel B focuses on portfolio characteristics: *Wealth*, the account balance; *NumAssets*, the number of assets held by the investor across accounts; *PctEquityShare*, the percentage of wealth in Equities—directly or through mutual funds; *PctBondShare*, the percentage of wealth in corporate bonds—directly or through mutual funds; *PctCashShare*, the percentage of wealth money market mutual funds—directly or through mutual funds; *PctMutualFunds*, the percentage of wealth directly invested in mutual funds; *PctCash*, the percentage of wealth directly invested in money market mutual funds; *PctStocks*, the percentage of wealth directly invested in individual stocks; *PctETF*, the percentage of wealth directly invested in ETFs; *PctIndex*, the percentage of mutual fund wealth invested in index funds; *PctInternational*, the percentage of mutual fund wealth invested in global or international funds—identified using the Lipper mutual fund classification; *PctEmerging*, the percentage of mutual fund wealth invested in emerging market funds—identified using the Lipper mutual fund classification. Panel C focuses on transactions and fees paid: *MgtFees*, the value-weighted management fees charged by the mutual funds held by the account-holders; *ExpRatio*, the value-weighted expense ratio charged by the mutual funds held by the investors; *TurnRatio*, the value-weighted turnover ratio of the mutual funds held by the investors; *Transaction*, the number of transactions directly initiated by the investors over the month before signing up for advice; *Volume*, is the volume (in US dollars) traded by the investors over the month before signing up for advice. For each variable, we report the number of accounts used in the computations, the mean, the standard deviation, and various percentiles of the distribution: the 10^{th}, 25^{th}, 50^{th}, 75^{th}, and 90^{th}.

Table 2. Demographic and Portfolio Characteristics of Advised Investors 12 Months after Sign-up.

Panel A. Demographic Characteristics

	N	mean	sd	p10	p25	p50	p75	p90
Age	54,328	63.81	12.39	46.00	57 .00	65.00	72.00	78.00
Male	54,747	0.59	0.49	0.00	0.00	1.00	1.00	1.00
Married	54,747	0.72	0.45	0.00	0.00	1.00	1.00	1.00
Tenure	54,747	13.55	9.30	1.92	3.83	13.66	20.25	25.75

Panel B. Portfolio Allocation

	N	mean	sd	p10	p25	p50	p75	p90
Wealth	54,747	758,340	821,023	108,923	210,784	478,890	981,327	1,764,745
NumAssets	54,747	7.95	4.91	4.00	5.00	6.00	9.00	14.00
PctEquityShare	54,742	0.59	0.18	0.39	0.49	0.57	0.69	0.85
PctBondShare	54,742	0.39	0.18	0.13	0.28	0.41	0.50	0.59
PctCashShare	54,742	0.02	0.05	0.00	0.00	0.00	0.01	0.05
PctMutualFunds	54,744	0.95	0.10	0.83	0.96	1.00	1.00	1.00
PctCash	54,744	0.02	0.05	0.00	0.00	0.00	0.01	0.05
PctStocks	54,744	0.01	0.04	0.00	0.00	0.00	0.00	0.03
PctETF	54,744	0.01	0.03	0.00	0.00	0.00	0.00	0.00
PctIndex	54,747	0.83	0.18	0.57	0.75	0.86	1.00	1.00
PctInternational	54,723	0.32	0.09	0.19	0.28	0.34	0.36	0.39
PctEmerging	54,723	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Panel C. Transactions and Fees

	N	mean	sd	p10	p25	p50	p75	p90
MgtFee	54,717	0.07	0.02	0.06	0.06	0.06	0.07	0.10
ExpRatio($\times$100)	54,707	0.09	0.03	0.07	0.08	0.08	0.10	0.13
TurnRatio	54,685	0.27	0.12	0.09	0.19	0.28	0.34	0.40
Transaction	54,747	2.63	4.09	0.00	0.00	1.00	3.00	8.00
Volume ($)	54,747	15,567	50,131	0.00	0.00	244	4,000	30,286

This table reports 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: Age, the age of the investor as of December 2017; $Male$, whether the investor is male; $Married$, whether he/she is married; $Tenure$, the tenure of the investor as of December 2017. Panel B focuses on portfolio characteristics: $Wealth$, the account balance; $NumAssets$, the number of assets held by the investor across accounts; *PctEquityShare*, the percentage of wealth in Equities—directly or through mutual funds; *PctBondShare*, the percentage of wealth in corporate bonds—directly or through mutual funds; $PctCashShare$, the percentage of wealth money market mutual funds—directly or through mutual funds; $PctMutualFunds$, the percentage of wealth directly invested in mutual funds; $PctCash$, the percentage of wealth directly invested in money market mutual funds; *PctStocks*, the percentage of wealth directly invested in individual stocks; $PctETF$, the percentage of wealth directly invested in ETFs; *PctIndex*, the percentage of mutual fund wealth invested in index funds; *PctInternational*, the percentage of mutual fund wealth invested in global or international funds—identified using the Lipper mutual fund classification; *PctEmerging*, the percentage of mutual fund wealth invested in emerging market funds—identified using the Lipper mutual fund classification. Panel C focuses on transactions and fees paid: $MgtFees$, the value-weighted management fees charged by the mutual funds held by the account-holders; *ExpRatio*, the value-weighted expense ratio charged by the mutual funds held by the investors; $TurnRatio$, the value-weighted turnover ratio of the mutual funds held by the investors; $Transaction$, the number of transactions directly initiated by the investors over the month before signing up for advice; $Volume$, the volume (in US dollars) traded by the investors over the month before signing up for advice. For each variable, we report the number of accounts used in the computations, the mean, the standard deviation, and various percentiles of the distribution: the 10^{th}, 25^{th}, 50^{th}, 75^{th}, and 90^{th}.

Table 3. Abnormal Sharpe Ratio at Robo-Advising Adoption.
Results Net of Robo-advising Fees.

Panel A. 3-Month Horizon

	All Accounts		Matched Accounts		
	After	**Before**	**After**	**Before**	**Difference**
Sharpe Ratio	-0.051***	-0.006*	-0.004	-0.014***	0.010
	(-12.87)	(-1.74)	(-0.88)	(-3.55)	(1.55)
N	73,677	59,758	55,336	55,336	55,336

Panel B. 6-Month Horizon

	All Accounts		Matched Accounts		
	After	**Before**	**After**	**Before**	**Difference**
Sharpe Ratio	0.002	0.000	0.053***	-0.073***	0.127***
	(0.61)	(0.08)	(14.43)	(-23.91)	(24.05)
N	67,787	52,093	43,202	43,202	43,202

Panel C. 9-Month Horizon

	All Accounts		Matched Accounts		
	After	**Before**	**After**	**Before**	**Difference**
Sharpe Ratio	0.012***	-0.019***	-0.001	0.004	-0.005
	(5.08)	(-7.39)	(-0.20)	(1.51)	(-0.97)
N	60,222	45,665	31,482	31,482	31,482

This table reports investment performance across investors before and after signing up for advice. The results are computed at the investor level. As a measure of performance, we use the abnormal Sharpe Ratio, that is, the difference between the realized Sharpe ratio of each account-holder and the realized Sharpe ratio of the market portfolio, where the latter is computed as the value-weighted returns on the NYSE/AMEX/NASDAQ CRSP portfolio. For each account-holder, portfolio returns and volatilities are computed using beginning-of-month investment holdings. Furthermore, portfolio volatilities are computed as realized volatilities using squared daily returns. Panels A, B, and C report results for Sharpe ratios computed at the 3-, 6-, and 9-month horizons. Each column within each panel reports average abnormal Sharpe ratios, t-statistics testing the null hypothesis that the averages are equal to zero, and the number of observations used in the computations of the results. The first and second column report statistics for all investors available before and after signing up for advice. The third and fourth column uses only the accounts present before and after. Finally, the last column reports column reports the average performance difference after and before signing up for advice, computed at the investor level. In all cases, performance is computed from the time the investor signs up for the service.

Table 4. Abnormal Sharpe Ratio at Six Months Before and After Robo-Advising Adoption. Results Net of Robo-advising Fees.

Panel A. 3-Month Horizon

	All Accounts		Matched Accounts		
	After	**Before**	**After**	**Before**	**Difference**
Sharpe Ratio	0.103***	-0.013***	0.104***	0.070***	0.034***
	(28.97)	(-3.23)	(19.15)	(19.14)	(5.26)
N	65,061	48,008	35,409	35,409	35,409

Panel B. 6-Month Horizon

	All Accounts		Matched Accounts		
	After	**Before**	**After**	**Before**	**Difference**
Sharpe Ratio	0.115***	-0.014***	0.125***	0.127***	-0.002
	(39.08)	(-4.68)	(22.13)	(38.43)	(-0.37)
N	56,232	41,366	22,929	22,929	22,929

Panel C. 9-Month Horizon

	All Accounts		Matched Accounts		
	After	**Before**	**After**	**Before**	**Difference**
Sharpe Ratio	0.094***	0.021***	0.432***	0.109***	0.323***
	(36.82)	(7.47)	(79.26)	(30.50)	(51.11)
N	47,839	35,024	11,252	11,252	11,252

This table reports investment performance across all investors before and after signing up for advice. The results are computed at the investor level. As a measure of performance, we use the abnormal Sharpe Ratio, that is, the difference between the realized Sharpe ratio of each account-holder and the realized Sharpe ratio of the market portfolio, where the latter is computed as the value-weighted returns on the NYSE/AMEX/NASDAQ CRSP portfolio. For each account-holder, portfolio returns and volatilities are computed using beginning-of-month investment holdings. Furthermore, portfolio volatilities are computed as realized volatilities using squared daily returns. Panels A, B, and C report results for Sharpe ratios computed at the 3-, 6-, and 9-month horizons. Each column within each panel reports average abnormal Sharpe ratios, t-statistics testing the null hypothesis that the averages are equal to zero, and the number of observations used in the computations of the results. The first and second column report statistics for all investors available before and after signing up for advice. The third and fourth column uses only the accounts present before and after. Finally, the last column reports column reports the average performance difference after and before signing up for advice, computed at the investor level. In all cases, performance is computed starting from 6 months before and after advice is implemented for the investor.

Table 5. Panel Regressions Relating Robo-Advising Adoption and Performance. Results Net of Robo-advising Fees.

Panel A. Sharpe Ratio

	Spec_1	Spec_2	Spec_3	Spec_4	Spec_5	Spec_6
ROBO	0.885***	0.885***	2.051***	2.051***	0.364***	0.364***
	(275.48)	(3.58)	(473.76)	(3.33)	(100.74)	(3.20)
Fixed Effects	✗	✗	✓	✓	✓	✓
Time Effects	✗	✗	✗	✗	✓	✓
SE Clustering	Single	Double	Single	Double	Single	Double

Panel B. Returns Results

	Spec_1	Spec_2	Spec_3	Spec_4	Spec_5	Spec_6
ROBO	0.023***	0.023	0.068***	0.068	-0.001***	-0.001
	(153.64)	(1.50)	(319.51)	(1.65)	(-4.12)	(-0.14)
Fixed Effects	✗	✗	✓	✓	✓	✓
Time Effects	✗	✗	✗	✗	✓	✓
SE Clustering	Single	Double	Single	Double	Single	Double

Panel C. Volatility Results

	Spec_1	Spec_2	Spec_3	Spec_4	Spec_5	Spec_6
ROBO	-0.017***	-0.017***	-0.035***	-0.035***	-0.009***	-0.009***
	(-197.07)	(-8.28)	(-275.10)	(-6.66)	(-70.98)	(-19.27)
Fixed Effects	✗	✗	✓	✓	✓	✓
Time Effects	✗	✗	✗	✗	✓	✓
SE Clustering	Single	Double	Single	Double	Single	Double

This table reports results relating robo-advising adoption and investment performance. The baseline regression estimated is:

$$Sharpe_{i,t} = \alpha_i + \beta_t + \gamma \; ROBO_{i,t} + \epsilon_{i,t},$$

where $Sharpe_{i,t}$ represents the monthly Sharpe ratio, α_i denotes individual fixed effects, β_t monthly time-effects, and $ROBO$ is an indicator variable equal to one when the portfolio allocation is performed by the robo-advisor and zero otherwise. In each panel, specifications 1 and 2 do not include individual fixed effects, while specifications 3 through 6 do. Specifications 1 through 4 do not include time-effects, while specifications 5 and 6 do. Finally, in specifications 1, 3, and 5, standard errors are clustered at the individual level. Standard errors are instead clustered double-clustered in specifications 2, 4, and 6. Panel A reports the results for portfolios' annualized Sharpe ratios. Panel B repeats the analysis for portfolio annualized returns. Finally, Panel C reports results for annualized portfolio volatility.

Table 6. Identification Results Using Matched Investor Returns as Benchmarks

	Top Decile	**Top 2 Deciles**	**Top 3 Deciles**	**All Investors**
Difference	-0.069***	-0.071***	-0.074***	-0.072***
	(-29.98)	(-30.74)	(-31.60)	(-30.96)
N	297,134	297,134	297,134	297,134

This table reports the results of an identification strategy procedure that controls for investors' endogenous decision to sign-up for robo-advice. The results are computed as follows. We first consider only those individuals who have signed-up for advice during the sample and focus on their returns starting from 6 months after they have signed-up for advice. We then construct, for each investor, the correlation of the monthly realized sharpe with every other investor that has been advised. For each investor, we rank all other investors in terms of their risk-return correlations when they are advised. Intuitively, investors that have the same risk-return profiles when advised are very similar in their characteristics—at least from the robo-advisor perspective. Finally, we exploit the fact that investors sign up at different times to construct counter-factual returns for each investors. We provide a simple example of the strategy below. As shown in the diagram, suppose we have three investors A, B and C and we want to construct the counterfactual performance for investor A. Suppose also that A signed up for advice on January 1st 2016, while B and C signed up for advice on July 1st 2014. Finally, assume that the three investors are advised through the end of 2019. We use the period starting from July 1st 2016 through December 2019—the are highlighted in green—to compute the portfolio risk-return correlation of investor A with both investors B and C. Suppose that the correlation with B is 92% while the correlation with C is 99%. For investor A, we would rank investor C as first and investor B as second in the construction of the counterfactual. In the final step, we would compare the performance of investor A during the period January 1st 2015 through December 31st 2015—when he/she was unadvised—to the performance of investor C over the same period—when he/she was already advised. In the diagram below, this is denoted using the red shaded area.

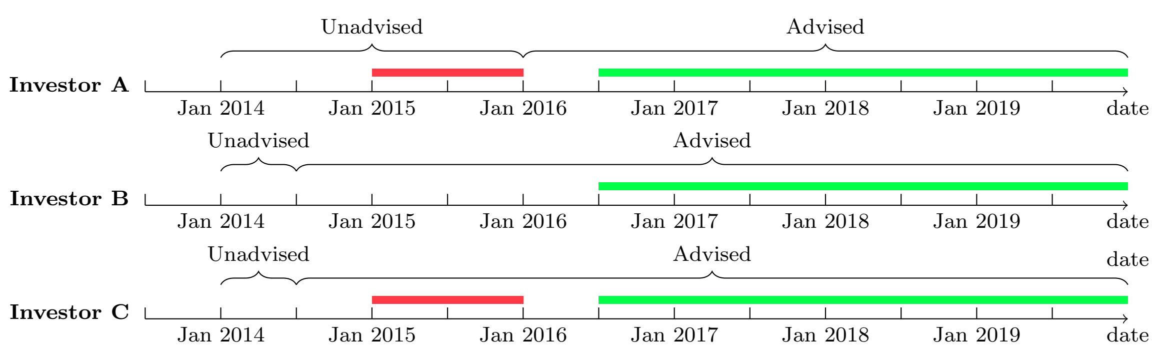

The first column compares the performance of the unadvised individuals to the top decile of closest matches. The second and third column focuses on the top 2 and 3 deciles, respectively. Finally, the fourth column focuses on the all matched investors.

Table 7. Cross-sectional Results on the Determinants of Robo-advising Adoption

	Linear Probability Model	**Logit Model**
Market_Ret	-0.001	0.000
	(-0.76)	(0.12)
Market_Var	0.023***	0.021***
	(14.54)	(13.96)
Investor_Ret	-0.005***	-0.006***
	(-2.83)	(-3.48)
Realized_Var	0.016***	0.016***
	(10.51)	(10.74)
Percent_ETF	0.014***	0.012***
	(6.65)	(5.73)
Percent_MF	0.002	-0.002
	(0.47)	(-0.36)
Percent_Stock	0.012***	0.010***
	(5.58)	(4.70)
Percent_International	-0.012***	-0.012***
	(-12.58)	(-12.02)
Expense_Ratio	0.051***	0.045***
	(49.22)	(44.67)
Client_Wealth	-0.037***	-0.036***
	(-36.32)	(-35.92)
Equity_Share	-0.029***	-0.029***
	(-17.66)	(-17.35)
Cash_Share	-0.005	-0.009**
	(-1.20)	(-2.00)
Human_Advisor_Scheduled_Appointments	0.076***	0.072***
	(48.49)	(47.64)
Investor_Scheduled_Appointments	0.004	0.003
	(1.30)	(0.95)
Constant	0.212***	
	(174.77)	
R-square	0.049	0.042
N	194,843	194,843

This table reports cross-sectional regression results on the determinants of robo-advising adoption. The baseline regression estimated is:

$$Sign\_Up_i = \alpha + \beta\, \boldsymbol{x}_i + \epsilon_i$$

where $Sign\_Up_i$ is an indicator variable equals to 0 if the investor does not adopt advice and 1 if it does. The vector $\boldsymbol{x}_i$ contains the following market-wide regressors: $Market\_Ret$ and $Market\_Var$ are market-wide returns and volatility, respectively. It also includes the following individual specific returns and portfolio-characteristic variables: $Investor\_Ret$, is the investor's monthly return; $Realized\_Var$ is the monthly realized variance; $Percent\_ETF$ is the client percentage of wealth in ETF; $Percent\_MF$ is the percentage of wealth in mutual funds; $Percent\_Stock$ is the percentage of wealth in individual stocks; $Expense\_ratio$, is the expense ratio paid on the mutual funds held; $Percent\_International$ is the percentage of wealth in international stocks or bonds; $Client\_wealth$ is overall client wealth of the investor; finally, $Equity\_share$ and $Cash\_share$ are the investors' portfolio shares in equities and cash (including money market mutual funds), respectively. The last two regressors included are $Human\_Advisor\_Scheduled\_Appointments$ and $Investor\_Scheduled\_Appointments$. These appointments include appointments categorized as client-facing, individual consults, plan delivery, ad-hoc meetings, check-in, and plan delivery. All non-binary regressors are standardized so that they have unit variance. This helps with the interpretation of the results. The first column reports results for a linear probability model. The second column repeats the analysis using a Logit specification and reports the marginal effects.

Table 8. Cox-Proportional Hazard Model Results on Robo-advising Attrition

	Spec 1	**Spec 2**
Market_Ret	0.990	1.004
	(-0.33)	(0.14)
Market_Var	0.907***	0.918***
	(-3.63)	(-3.17)
Investor_Ret	0.982	0.976
	(-0.63)	(-0.83)
Realized_Var	1.048*	1.044*
	(1.85)	(1.68)
Percent_ETF	1.044	1.050
	(1.33)	(1.48)
Percent_MF	0.929	0.947
	(-0.89)	(-0.65)
Percent_Stock	1.012	1.012
	(0.33)	(0.33)
Percent_International	0.995	1.001
	(-0.30)	(0.07)
Expense_Ratio	0.910***	0.899***
	(-4.48)	(-5.04)
Client_Wealth	0.984	1.220***
	(-0.88)	(7.66)
Equity_Share	1.104***	1.119***
	(3.13)	(3.52)
Cash_Share	1.080	1.095
	(1.02)	(1.18)
Human_Advisor_Scheduled_Appointments	1.028*	1.027
	(1.66)	(1.62)
Investor_Scheduled_Appointments	0.947***	0.963**
	(-2.88)	(-1.98)
"Level 1" Dummy		2.084***
		(14.02)
"Level 2" Dummy		1.408***
		(6.92)
N	46,862	46,862

This table reports cox proportional hazard results on the determinants of attrition from robo-advising. The hazard function takes the following form:

$$\lambda(t|\boldsymbol{x}_i) = \lambda_0(t)\exp(\boldsymbol{x}_i \cdot \beta),$$

where $\lambda(t|\boldsymbol{x}_i)$ is the hazard function at time t for investor i with covariate vector (explanatory variables) $\boldsymbol{x}_i$. The vector $\boldsymbol{x}_i$ contains the following market-wide regressors: $Market\_Ret$ and $Market\_Var$ are market-wide returns and volatility, respectively. It also includes the following individual specific returns and portfolio-characteristic variables: $Investor\_Ret$, is the investor's monthly return; $Realized\_Var$ is the monthly realized variance; $Percent\_ETF$ is the client percentage of wealth in ETF; $Percent\_MF$ is the percentage of wealth in mutual funds; $Percent\_Stock$ is the percentage of wealth in individual stocks; $Expense\_ratio$, is the expense ratio paid on the mutual funds held; $Percent\_International$ is the percentage of wealth in international stocks or bonds; $Client\_wealth$ is overall client wealth of the investor; finally, $Equity\_share$ and $Cash\_share$ are the investors' portfolio shares in equities and cash (including money market mutual funds), respectively. The last two regressors included are $Human\_Advisor\_Scheduled\_Appointments$ and $Investor\_Scheduled\_Appointments$. These appointments include appointments categorized as client-facing, individual consults, plan delivery, ad-hoc meetings, check-in, and plan delivery. Reported are hazard ratios rather than coefficient estimates to ease the interpretation of the results. The first column reports the baseline results. The second column includes the two client segmentation dummies "Level 1" and "Level 2" investors. Level 1 have \$500K at most in AUM, while Level 2 investors have between \$500K and \$1M in AUM.

Online Appendix for the Paper:

Who Benefits from Robo-Advising? Evidence from Machine Learning

Alberto G. Rossi and Stephen Utkus

(Not for publication)

Online Appendix A.1 Holdings of Advised and non-Advised Investors

To better understand how the service manages the wealth of its investors, we compare the main holdings of advised and non-advised investors. Overall, we find advice lowers the number of assets held across investors. We show it in three way. First, we present the top tickers held across advised and non-advised investors in January 2017.[8] These are reported in Table Online I. Starting with mutual funds in Panel A, advised investors are invested in four mutual funds, i.e. VTSAX, VTIAX, VBTLX, and VTABX. Combined, these four mutual funds represent almost 75% of advised wealth in January 2017. The mix of funds is designed to expose investors to both US and international diversified stock and bond portfolios.

VTSAX is the Vanguard Total Stock Market Index Fund Admiral Shares. It is the indexed mutual fund equivalent of the VTI ETF. Its benchmark is the CRSP US Total Market Index, which represents 100% of the CRSP US stock market index. The fees are extremely low: management fees are 0.03% per year and total annual operating expenses of only 0.04%. VTIAX is the Vanguard Total International Stock Index Fund Admiral Shares. The fund invests in European Equities (42%), Pacific Region (30%), Emerging Markets (21%) and North America (6.6%). The fees of this fund are also rather low, 0.08% management fees and 0.11% total annual operating expenses. The last two are bond funds. VBTLX is the Vanguard Total Bond Market Index Fund Admiral Shares and invests in public, investment-grade, taxable, fixed income securities in the US—including government bonds, corporate bonds, mortgage-backed and asset-backed securities. The fees are, once again, extremely low: 0.04% management fees and 0.05% total annual operating expenses. Finally, VTABX is the Vanguard Total International Bond Index Fund Admiral Shares. The fund invests in government, government agency, corporate, and securitized non-U.S. investment-grade fixed income investments, all issued in currencies other than the U.S. dollar and with maturities of more than one year. The fund features 0.09% management fees and 0.11% total annual operating expenses.

We see a similar effect in ETFs, where 30% of the advised investor wealth is invested in VTI. In terms of stock holdings, we do not see any effect. Finally, turning to cash, we find that all the wealth of advised investors is concentrated in the money market funds VMMXX, VMSXX, and VUSXX. Together, the three comprise 89% of investors' wealth.

[8] The results are very similar if we compute the results for other dates.

Next, we analyze whether the concentration of wealth across tickers has varied over time. We also measure the extent to which wealth is dispersed across few or many assets in the cross-section. The results for the first exercise is reported in Figure Online III, where we report—for each asset class—the proportion of wealth allocated to the top 5 tickers— across all Vanguard investors. The red dotted line represent results for advised investors while the black solid line the results for non-advised investors. The results in Figure Online III are very much in line with those in Table Online I. The cross-sectional results are reported in Figure Online IV in that portfolios of advised investors are much more consolidated than non-advised investors. The figure plots the cumulative wealth invested across tickers, from the more to the least purchased. The red dotted line represent results for advised investors while the black solid line the results for non-advised investors. The vertical red line represents the total number of either mutual funds, ETF, stocks, or money market funds held by each investor category, that is, advised and non-advised. The results show advised investors are invested in approximately 3,000 mutual funds in total, while non-advised investors invest in almost 7,000 funds, that is, wealth is concentrated in fewer funds for advised investors, compared to non-advised investors. Note that this does not entail that investors are not well-diversified as mutual funds themselves hold broadly diversified portfolios. Also note that the large cross-section of mutual funds held by advised investors is likely due to investors' willingness to maintain in their portfolio specific mutual funds that are not managed directly by the robo-advisor.[9] Largely, the results for ETFs and money market mutual funds are similar. Advised investors have more concentrated portfolios than non-advised investors. The total number of ETFs held across non-advised investors is close to 1300, while the one for advised investors is less than 650. Finally, advised investors also invest in less individual stocks. Non-advised investors collectively invest in as many as 8,500 individual stocks. The corresponding number of stocks for advised investors is a little more than 2,000.

[9]Investors can choose to hold specific investment vehicles not managed by the robo-advisor, even when they are advised.

Online Appendix B.1 Boosted Regression Trees

This online appendix provides a brief introduction to Boosted Regression Trees. Section Online Appendix B.1.1 describes Regression Trees, Section Online Appendix B.1.2 describes Boosting. Finally, Section Online Appendix B.1.3 describes the implementation of BRT adopted in the paper.[10]

Online Appendix B.1.1 Regression Trees

Suppose we have P potential predictor ("state") variables and a single dependent variable over T observations, i.e. (x_t, y_{t+1}) for $t = 1, 2, ..., T$, with $x_t = (x_{t1}, x_{t2}, ..., x_{tp})$. Fitting a regression tree requires deciding (i) which predictor variables to use to split the sample space and (ii) which split points to use. The regression trees we use employ recursive binary partitions, so the fit of a regression tree can be written as an additive model:

$$f(x) = \sum_{j=1}^{J} c_j I\{x \in S_j\}, \tag{4}$$

where $S_j,\ j = 1, ..., J$ are the regions we split the space spanned by the predictor variables into, $I\{\}$ is an indicator variable and c_j is the constant used to model the dependent variable in each region. If the L^2 norm criterion function is adopted, the optimal constant is $\hat{c}_j = mean(y_{t+1}|x_t \in S_j)$.

The globally optimal splitting point is difficult to determine, particularly in cases where the number of state variables is large. Hence, a sequential greedy algorithm is employed. Using the full set of data, the algorithm considers a splitting variable p and a split point s so as to construct half-planes

$$S_1(p, s) = \{X|X_p \leq s\} \qquad and \qquad S_2(p, s) = \{X|X_p > s\}$$

that minimize the sum of squared residuals:

$$\min_{p,s} \left[\min_{c_1} \sum_{x_t \in S_1(p,s)} (y_{t+1} - c_1)^2 + \min_{c_2} \sum_{x_t \in S_2(p,s)} (y_{t+1} - c_2)^2 \right]. \tag{5}$$

[10]Our description draws on Friedman, Hastie, and Tibshirani (2001), who provide a more in-depth coverage of the approach.

For a given choice of p and s the fitted values, $\hat{c}_1$ and $\hat{c}_2$, are

$$\begin{aligned}\widehat{c}_1 &= \frac{1}{\sum_{t=1}^{T} I\{x_t \in S_1(p,s)\}} \sum_{t=1}^{T} y_{t+1} I\{x_t \in S_1(p,s)\}, \\ \widehat{c}_2 &= \frac{1}{\sum_{t=1}^{T} I\{x_t \in S_2(p,s)\}} \sum_{t=1}^{T} y_{t+1} I\{x_t \in S_2(p,s)\}. \end{aligned} \tag{6}$$

The best splitting pair (p, s) in the first iteration can be determined by searching through each of the predictor variables, $p = 1, .., P$. Given the best partition from the first step, the data is then partitioned into two additional states and the splitting process is repeated for each of the subsequent partitions. Predictor variables that are never used to split the sample space do not influence the fit of the model, so the choice of splitting variable effectively performs variable selection.

Regression trees are generally employed in high-dimensional datasets where the relation between predictor and predicted variables is potentially non-linear. This becomes important in our context as it is not clear which variables may be more or less relevant *ex-ante*. Furthermore, it is difficult to know in our context whether there is a linear relation between predictor and predicted variables. On the other hand, the approach is sequential and successive splits are performed on fewer and fewer observations, increasing the risk of fitting idiosyncratic data patterns. Furthermore, there is no guarantee that the sequential splitting algorithm leads to the globally optimal solution. To deal with these problems, we next consider a method known as boosting.

Online Appendix B.1.2 Boosting

Boosting is based on the idea that combining a series of simple prediction models can lead to more accurate forecasts than those available from any individual model. Boosting algorithms iteratively re-weight data used in the initial fit by adding new trees in a way that increases the weight on observations modeled poorly by the existing collection of trees. From above, recall that a regression tree can be written as:

$$\mathcal{T}\left(x; \{S_j, c_j\}_{j=1}^{J}\right) = \sum_{j=1}^{J} c_j I\{x \in S_j\} \tag{7}$$

A boosted regression tree is simply the sum of regression trees:

$$f_B(x) = \sum_{b=1}^{B} \mathcal{T}_b\left(x; \{S_{b,j}, c_{b,j}\}_{j=1}^{J}\right), \tag{8}$$

where $\mathcal{T}_b\left(x; \{S_{b,j}, c_{b,j}\}_{j=1}^{J}\right)$ is the regression tree used in the b-th boosting iteration and B is the number of boosting iterations. Given the model fitted up to the $(b-1)-th$ boosting iteration, $f_{b-1}(x)$, the subsequent boosting iteration seeks to find parameters $\{S_{j,b}, c_{j,b}\}_{j=1}^{J}$ for the next tree to solve a problem of the form

$$\{\hat{S}_{j,b}, \hat{c}_{j,b}\}_{j=1}^{J} = \min_{\{S_{j,b}, c_{j,b}\}_{j=1}^{J}} \sum_{t=0}^{T-1} \left[y_{t+1} - \left(f_{b-1}(x_t) + \mathcal{T}_b\left(x_t; \{S_{j,b}, c_{j,b}\}_{j=1}^{J}\right)\right)\right]^2. \tag{9}$$

For a given set of state definitions ("splits"), $S_{j,b}$, $j = 1, .., J$, the optimal constants, $c_{j,b}$, in each state are derived iteratively from the solution to the problem

$$\begin{aligned} \hat{c}_{j,b} &= \min_{c_{j,b}} \sum_{x_t \in S_{j,b}} \left[y_{t+1} - (f_{b-1}(x_t) + c_{j,b})\right]^2 \\ &= \min_{c_{j,b}} \sum_{x_t \in S_{j,b}} \left[e_{t+1,b-1} - c_{j,b}\right]^2, \end{aligned} \tag{10}$$

where $e_{t+1,b-1} = y_{t+1} - f_{b-1}(x_t)$ is the empirical error after $b-1$ boosting iterations. The solution to this is the regression tree that most reduces the average of the squared residuals $\sum_{t=1}^{T} e_{t+1,b-1}^2$ and $\hat{c}_{j,b}$ is the mean of the residuals in the jth state.

Forecasts are simple to generate from this approach. The boosted regression tree is first estimated using data from $t = 1, ..., t^*$. Then the forecast of y_{t^*+1} is based on the model estimates and the value of the predictor variable at time t^*, x_{t^*}. Boosting makes it more attractive to employ small trees (characterized by only two terminal nodes) at each boosting iteration, reducing the risk that the regression trees will overfit. Moreover, by summing over a sequence of trees, boosting performs a type of model averaging that increases the stability and accuracy of the forecasts.[11]

[11] See Rapach, Strauss, and Zhou (2010) for similar results in the context of linear regression.

Online Appendix B.1.3 Implementation

Our estimations follow the stochastic gradient boosting approach of Friedman (2001) and Friedman (2002) with $J = 2$ nodes. The baseline implementation employs 10,000 boosting iterations, but we conduct a number of robustness checks to show that the results are not very sensitive to this choice.

We adopt two refinements to the basic boosted regression tree methodology. The first is *shrinkage*. As with ridge regression and neural networks, shrinkage is a simple regularization technique that diminishes the risk of over-fitting by slowing the rate at which the empirical risk is minimized on the training sample. We use a shrinkage parameter, $0 < \lambda < 1$, which determines how much each boosting iteration contributes to the overall fit:

$$f_b(x) = f_{b-1}(x) + \lambda \sum_{j=1}^{J} c_{j,b} I\{x \in S_{j,b}\}. \tag{11}$$

Following common practice we set $\lambda = 0.001$ as it has been found (Friedman (2001)) that the best empirical strategy is to set λ very small and correspondingly increase the number of boosting iterations.

The second refinement is *subsampling* and is inspired by "bootstrap aggregation" (bagging), see Breiman (1996). Bagging is a technique that computes forecasts over bootstrap samples of the data and averages them in a second step, therefore reducing the variance of the final predictions. In our context, the procedure is adapted as follows: at each boosting iteration we sample without replacement one half of the training sample and fit the next tree on the sub-sample obtained.

Online Appendix B.1.4 Relative Influence Measures and Partial Dependence Plots

One criticism of machine learning algorithms is that they are "Black Boxes" that do not provide a lot of intuition to the researcher and the reader. This criticism is hardly applicable to Boosted Regression Trees that instead feature very useful and intuitive visualization tools.

Online Appendix B.1.4.1 Relative Influence measures. The first measure commonly used is generally referred to as "relative influence" measures. Consider the reduction in the empirical error every time one of the covariates $x_{\cdot,l}$, is used to split the tree. Summing the reductions in empirical errors (or improvements in fit) across the nodes in the tree gives a measure of the variable's influence

(Breiman et al. (1984)):

$$I_l(\mathcal{T}) = \sum_{j=2}^{J} \Delta e(j)^2 I(x(j) = l), \tag{12}$$

where $\Delta e(j)^2 = T^{-1} \sum_{t=1}^{T} \left(e_t(j-1)^2 - e_t(j)^2\right)$, is the reduction in the squared empirical error at the j'th node and $x(j)$ is the regressor chosen at this node, so $I(x(j) = l)$ equals one if regressor l is chosen and zero otherwise. The sum is computed across all observations, $t = 1, ..., T$ and over the $J - 1$ internal nodes of the tree.

The rationale for this measure is that at each node, one of the regressors gets selected to partition the sample space into two sub-states. The particular regressor at node j achieves the greatest reduction in the empirical risk of the model fitted up to node $j - 1$. The importance of each regressor, $x_{l,\cdot}$, is the sum of the reductions in the empirical errors computed over all internal nodes for which it was chosen as the splitting variable. If a regressor never gets chosen to conduct the splits, its influence is zero. Conversely, the more frequently a lag is used for splitting and the bigger its effect on reducing the model's empirical risk, the larger its influence.

This measure of influence can be generalized by averaging over the number of boosting iterations, B, which generally provides a more reliable measure of influence:

$$\bar{I}_l = \frac{1}{B} \sum_{b=1}^{B} I_l(\mathcal{T}_b). \tag{13}$$

This is best interpreted as a measure of relative influence that can be compared across regressors. We therefore report the following measure of relative influence, $\overline{RI}_l$, which sums to one:

$$\overline{RI}_l = \bar{I}_l / \sum_{l=1}^{L} \bar{I}_l. \tag{14}$$

Online Appendix B.1.4.2 Partial Dependence Plots. The second visualization tool featured by BRT are *partial dependence plots*, that are defined as follows. Suppose we select a particular covariate, X_p, from the set of P predictor variables $X = (X_1, X_2, ..., X_P)$ and denote the remaining variables X_{-p}, i.e. $X_{-p} = X \backslash \{X_p\}$. We use the following measure of the average marginal effect of X_p on the dependent variable

$$f_p(X_p) = E_{X_{-p}} f(X_p, X_{-p}). \tag{15}$$

This is called the average partial dependence measure. It fixes the value of X_p and averages out the effect of all other variables. By, repeating this process for different values of X_p, we trace out the marginal effect this covariate has on the predicted variable.

An estimate of $f_p(X_p)$ can be computed by averaging over the sample observations

$$\bar{f}_p(X_p) = \frac{1}{T}\sum_{t=1}^{T} f(X_p, x_{t,-p}), \tag{16}$$

where $x_{t,-p} = \{x_{1,-p}, ..., x_{T,-p}\}$ are the values of X_{-p} occurring in the data.

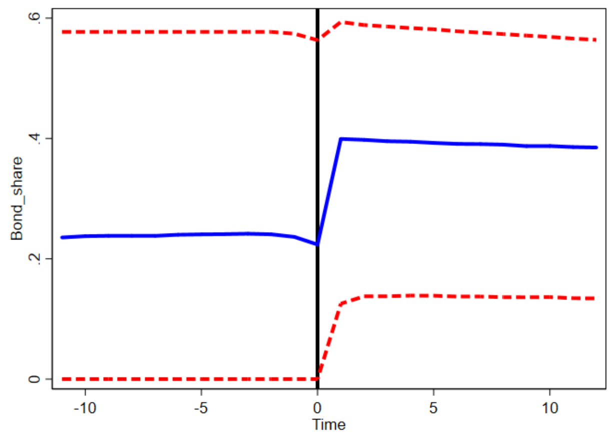

(a) Percentage of Bonds

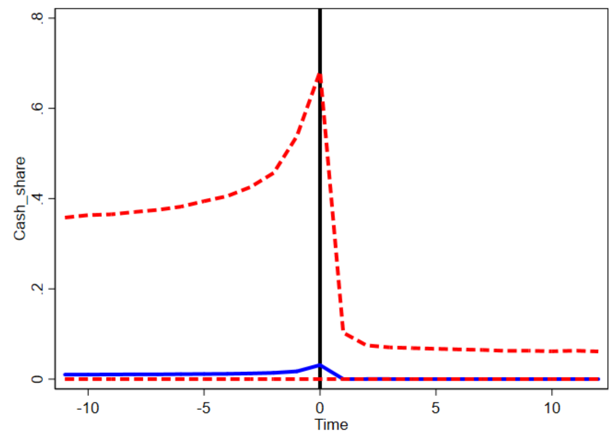

(b) Percentage of Cash

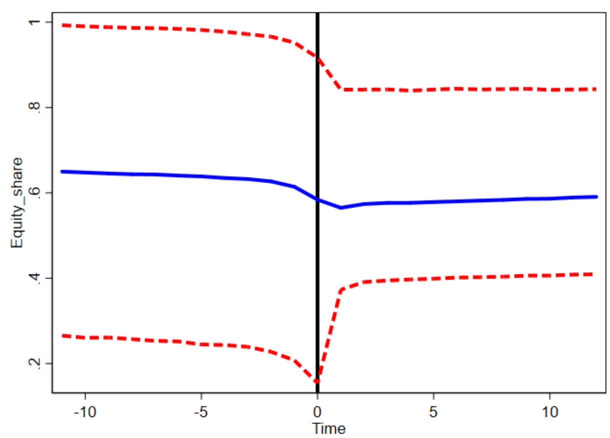

(c) Percentage of Equities

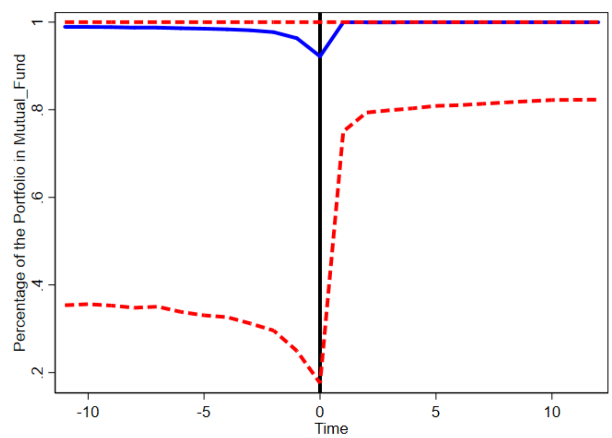

(d) Percentage in Mutual Funds

Figure Online I. Portfolio Allocation before and after advice: All Investors. This figure reports results for investor portfolio characteristics before and after signing up for robo-advice. The results are computed at the investor level and include all account types, that is, taxable and non-taxable (IRA) accounts. Subfigures (a), (b), (c) and (d) report results for the percentage of wealth held—directly or through mutual funds—in bonds, cash (including money market mutual funds), equities and mutual funds. In each subfigure, time "0" represent the month before investors sign up for advice. Results are computed using only investors that are in the sample for at least twelve months before and after signing up for advice. The blue line denotes median values, while the red dashed lines are the 10th and 90th percentiles of the distribution.

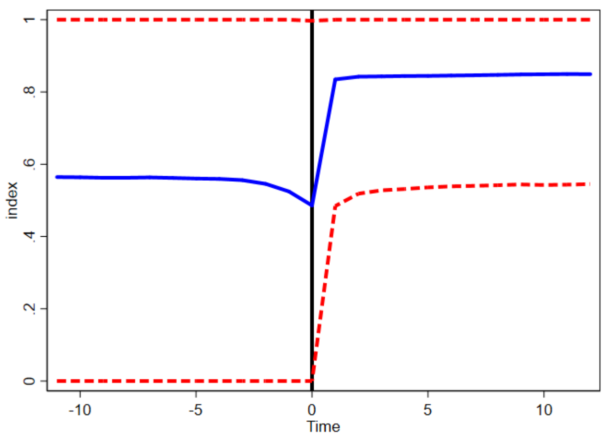

(a) Percentage of Indexed Funds

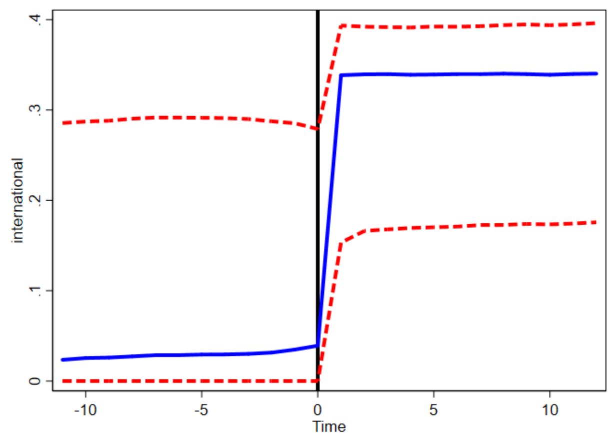

(b) Percentage of International Funds

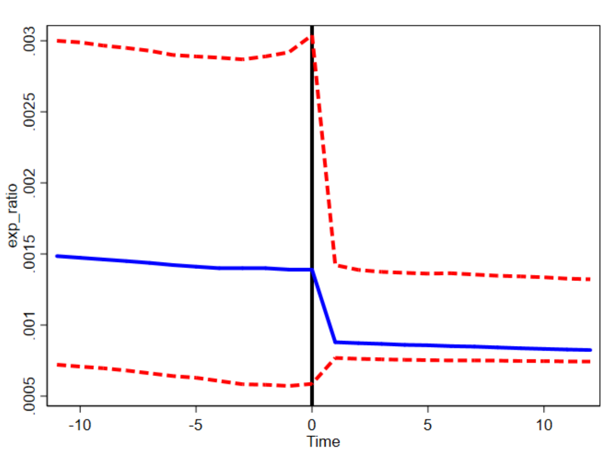

(c) Expense Ratios

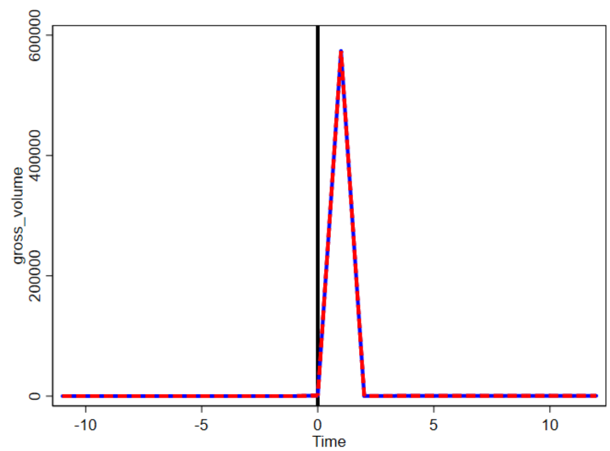

(d) Trading Volume

Figure Online II. Indexation, international diversification and fees before and after advice: All Investors. This figure reports results for investor portfolio characteristics before and after signing up for robo-advice. The results are computed at the investor level and include all account types, that is, taxable and non-taxable (IRA) accounts. Subfigure (a) reports results for the percentage of mutual fund wealth invested in indexed funds; Subfigure (b) the percentage of mutual fund wealth invested in global or international funds—identified using the Lipper mutual fund classification; Subfigure (c), the value-weighted expense ratio charged by the mutual funds held by the account-holders. Finally, Subfigure (d) shows results for the monthly trading volume, in US dollars. In each subfigure, time "0" represent the month before investors sign up for advice. Results are computed using only investors that are in the sample for at least twelve months before and after signing up for advice. The blue line denotes median values, while the red dashed lines are the 10th and 90th percentiles of the distribution.

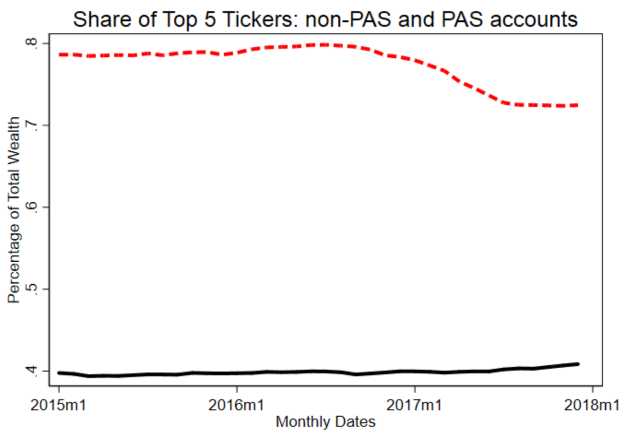

(a) Mutual Funds

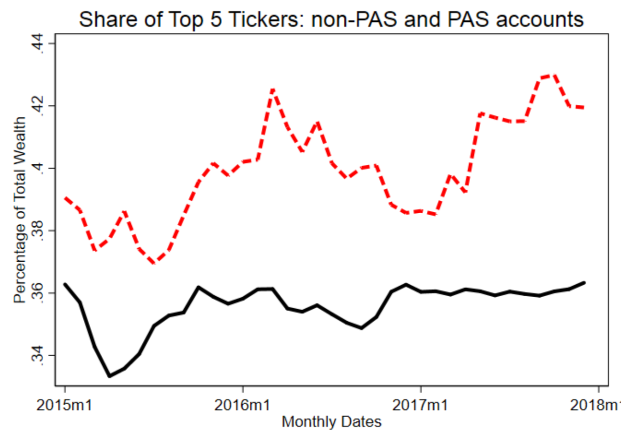

(b) ETF

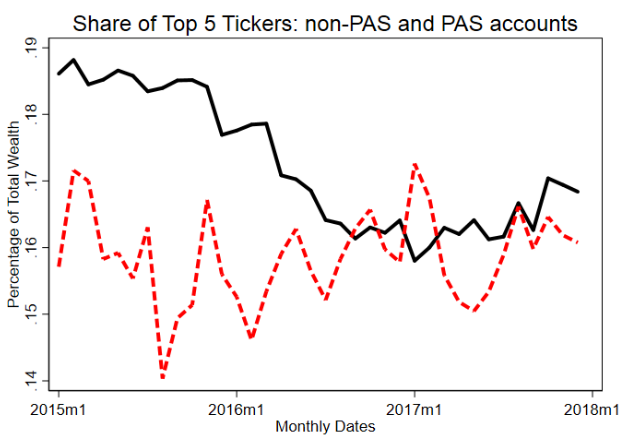

(c) Stocks

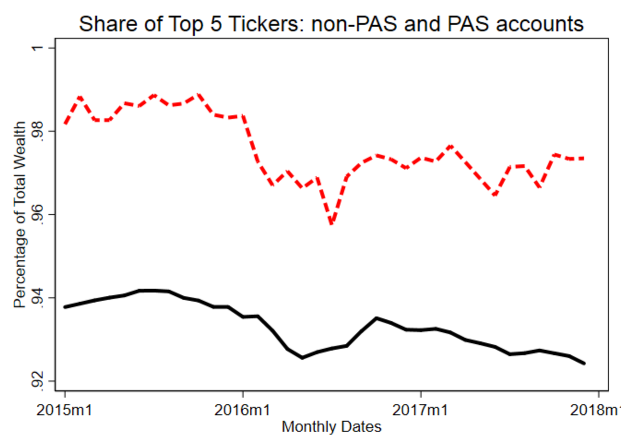

(d) Cash

Figure Online III. Percentage of Wealth in Top 5 Tickers by Asset Class Over Time. This figure reports the percentage of wealth in top 5 tickers over time, computed for non-advised and advised investors. The results are computed at the account-holder level and include all account types, that is, taxable and non-taxable (IRA) accounts. Subfigure (a) reports the results for mutual funds. Subfigure (b), (c), and (d) reports results for ETFs, individual stocks and money market mutual funds, respectively. In each subfigure, the black solid line presents the results for non-advised account holders, while the red dashed-line presents the results for advised account holders.

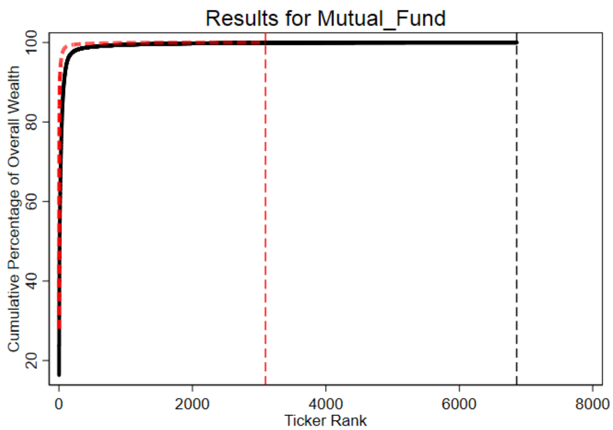

(a) Mutual Funds

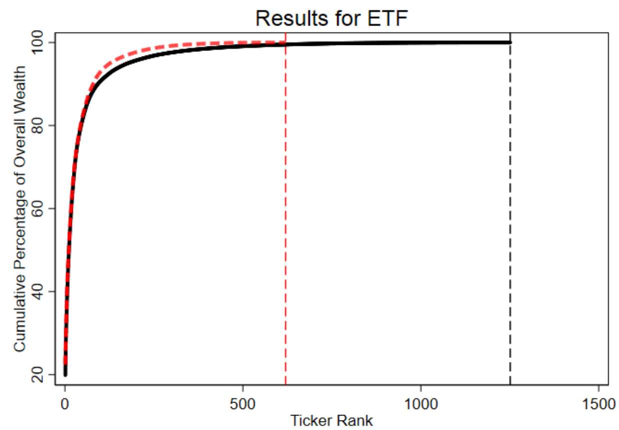

(b) ETF

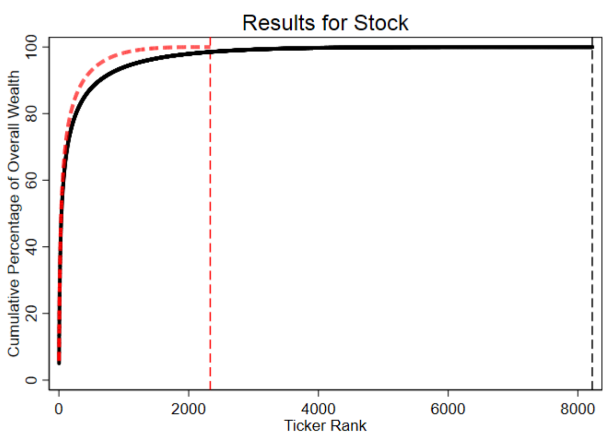

(c) Stocks

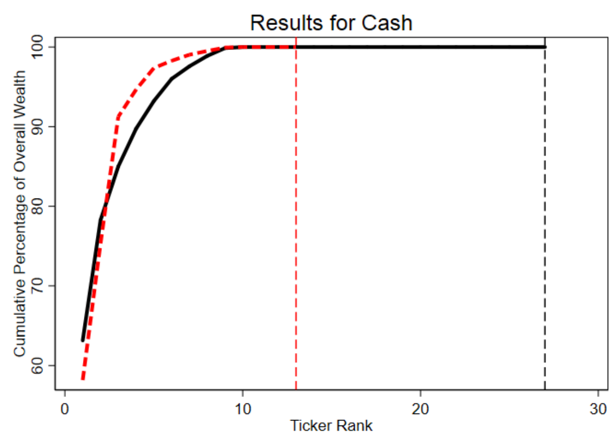

(d) Cash

Figure Online IV. Cumulative Percentage of Wealth Across Tickers in the Cross-section. This figure reports the cumulative percentage of wealth across all the tickers held in each asset class in January 2017, across non-advised and advised investors. The results are computed at the account-holder level and include all account types, that is, taxable and non-taxable (IRA) accounts. Subfigure (a) reports the results for mutual funds. Subfigure (b), (c), and (d) reports results for ETFs, individual stocks and money market mutual funds, respectively. In each subfigure, the black solid line presents the results for non-advised account holders, while the red dashed-line presents the results for advised account holders.

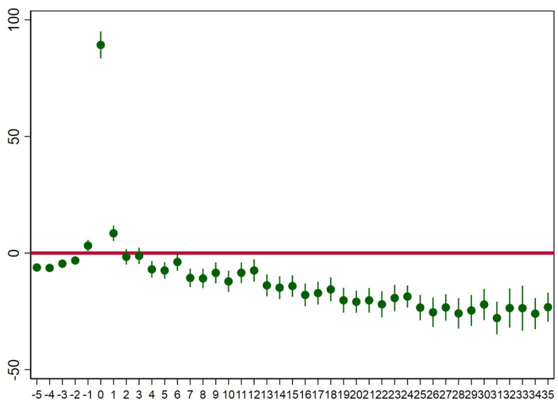

(a) Total

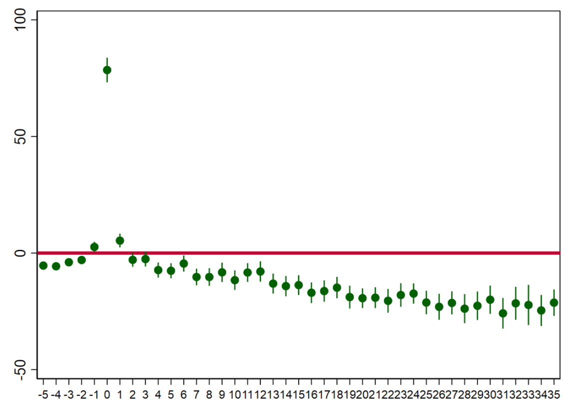

(b) Attention Through Desktop Computer

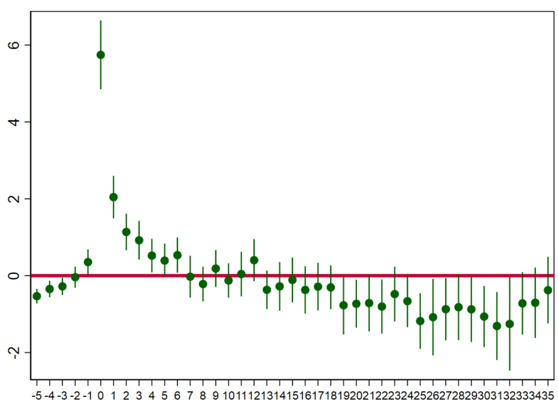

(c) Attention Through Mobile App

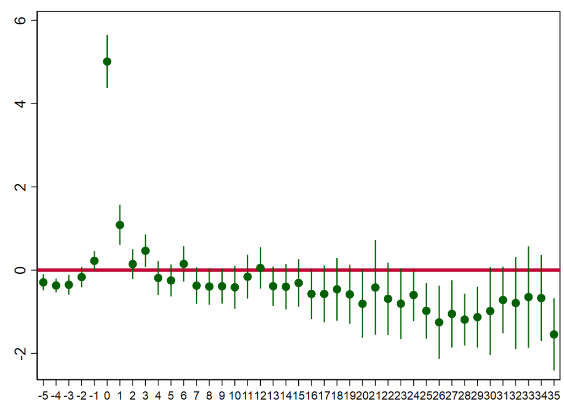

(d) Attention Through Mobile Browser

Figure Online V. Panel Regressions Relating Robo-Advising Adoption and Investors' Attention. This figure reports results relating robo-advising adoption and investment performance. The regression estimated is:

$$Attention_{i,t} = \alpha_i + \beta_t + \sum_{j=-5}^{35} \gamma_j \ ROBO_{i,j,t} + \delta \ X_{i,t} + \epsilon_{i,t},$$

where $Attention_{i,t}$ denotes the number of minutes spent on the robo-advisor website by investor i on month t, α_i denotes individual fixed effects, β_t monthly time-effects, and $ROBO_{i,j,t}$ is a dummy variable equal to zero for every month, except for the j-th month before robo-advising enrollment is initiated and after robo-advising is implemented. The $0-th$ month is the one when the robo-advisor is implemented, negative values of j refer to the months before the process of robo-advising is initiated. We stop at the $35-th$ month post-adoption as it is the longest horizon for which we have robo-advised clients. Note that in these results we do not include dummies for the periods of robo-advising enrollment as the enrollment process range between one to several months across individuals. The regressor $X_{i,t}$ contains investors' portfolio return and risk, which do not have an economically significant effect on attention. The figure is composed of four subfigures. Subfigure (a) reports the results for the total number of logins; subfigures (b) through (d) report results for the logins through desktop computers, mobile apps and mobile browsers, respectively. Each subfigure plots the γ_j coefficients and their 95% confidence intervals, computed using double-clustered standard errors.

Table Online I. Top Tickers by Asset Class in January 2017

Panel A. Mutual Fund

Rank	NON-Advised Ticker	NON-Advised Pct of Assets	Advised Ticker	Advised Pct of Assets
1	VTSAX	16%	VTSAX	28%
2	VFIAX	8%	VTIAX	18%
3	VBTLX	7%	VBTLX	15%
4	VTIAX	5%	VTABX	11%
5	VWIUX	4%	VFIDX	6%
6	VWENX	3%	VFSUX	4%
7	VGHAX	2%	VWIUX	3%
8	VWIAX	2%	VFIAX	2%
9	VTABX	2%	VMLUX	2%
10	VFSUX	2%	VEXAX	1%

Panel B. ETF

Rank	NON-Advised Ticker	NON-Advised Pct of Assets	Advised Ticker	Advised Pct of Assets
1	VTI	20%	VTI	22%
2	VOO	7%	SPY	5%
3	VYM	4%	VOO	4%
4	VIG	3%	VIG	3%
5	VXUS	3%	VXUS	3%
6	SPY	3%	BND	3%
7	VWO	3%	VUG	2%
8	VNQ	2%	IVV	2%
9	BND	2%	VYM	2%
10	VEU	2%	IWF	2%

Panel C. Stocks

Rank	NON-Advised Ticker	NON-Advised Pct of Assets	Advised Ticker	Advised Pct of Assets
1	AAPL	5%	AAPL	6%
2	XOM	3%	BRK B	4%
3	BRK B	3%	MSFT	3%
4	GE	2%	GE	3%
5	T	2%	XOM	2%
6	CVX	2%	JNJ	2%
7	JNJ	2%	T	2%
8	MSFT	2%	GOOGL	2%
9	BRK A	2%	CVX	2%
10	PG	1%	FB	1%

Panel D. Cash

Rank	NON-Advised Ticker	NON-Advised Pct of Assets	Advised Ticker	Advised Pct of Assets
1	VMMXX	63%	VMMXX	58%
2	VMSXX	15%	VMSXX	17%
3	VMFXX	7%	VMFXX	16%
4	VUSXX	5%	VMRXX	3%
5	VMRXX	3%	VCTXX	3%
6	VCTXX	3%	VUSXX	1%
7	VPTXX	2%	VPTXX	1%
8	VYFXX	1%	VYFXX	0%
9	VNJXX	1%	VNJXX	0%
10	VOHXX	0%	VOHXX	0%

This table reports the top 10 tickers held across NON-advised and advised investors. The results are computed at the account-holder level and include all account types, that is, taxable and non-taxable (IRA) accounts. Panel A reports the results for mutual funds. Panels B, C, and D reports results for ETFs, individual stocks and money market mutual funds, respectively. Within each panel, the left sub-panels report results for non-advised investors, while the right sub-panels report results for advised investors. Each sub-panel reports—from left to right—the holdings rank, the ticker, and the percentage of asset class wealth invested in the ticker, computed across account holders. The results are computed as of January 2017.