Journal of Financial Economics 157 (2024) 103869

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/finec

The diversification and welfare effects of robo-advising☆

Alberto G. Rossi[a,*], Stephen Utkus[a,b]

[a] Georgetown University, USA
[b] University of Pennsylvania, USA

ARTICLE INFO

Dataset link: Replication Package for "The Diversification and Welfare Effects of Robo-advising" (Reference data)

JEL classification:
D14
G11
O33

Keywords:
FinTech
Portfolio choice
Machine learning
Individual investors
Financial literacy
Technology adoption

ABSTRACT

We study the diversification and welfare effects of a large US robo-advisor on the portfolios of previously self-directed investors and document five facts. First, robo-advice reshapes portfolios by increasing indexing and reducing home bias, number of assets held, and fees. Second, these portfolio changes contribute to higher Sharpe ratios. Third, those who benefit most from robo-advice are investors who did not have high exposure to equities or indexing and had poorer diversification levels. Fourth, robo-advice decreases the time investors dedicate to managing their investments. Fifth, those investors who benefit most are more likely to join the service and not quit it.

Robo-advisors have surged in popularity recently 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 (Capponi et al., 2022). 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-advisors have the potential to be superior to human financial advisors, as the latter have been shown to display behavioral biases and cognitive limitations. As a result, robo-advisors are quickly attracting attention from policymakers and investors at all levels.

This paper provides a comprehensive analysis of the largest robo-advisor in the US—Personal Advisor Service (PAS) from Vanguard. As

☆ Toni Whited was the editor for this article and provided very valuable feedback throughout the process. The paper has benefited from the comments made at presentations at the NBER Big Data and Securities Markets meeting, NBER Summer Institute on Aging, AFA Annual Meeting, EFA Annual Meeting, New York Fed FinTech Conference, GFLEC Seminar, Virtual Finance Seminar Series in Corporate Finance and Investments, INFORMS Annual Meeting, ABFER 7th Annual Conference, Luohan Academy, 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, Queen Mary University, University of Technology Sydney, 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, Cyndy Pagliaro, Lin Peng, Alessandro Previtero, Tarun Ramadorai (the NBER discussant), Michael Reher (the AFA discussant), 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. Stephen Utkus was employed at Vanguard in a research capacity and retired in August 2020. Rossi is an unpaid consultant at Vanguard to access the anonymized data.
* Correspondence to: McDonough School of Business, Georgetown University, Washington, DC, USA.
E-mail address: agr60@georgetown.edu (A.G. Rossi).

https://doi.org/10.1016/j.jfineco.2024.103869
Received 8 June 2022; Received in revised form 9 May 2024; Accepted 12 May 2024
Available online 25 May 2024
0304-405X/

of February 2019, according to the website roboadvisorpros.com, PAS is the largest robo-advisor in the world with $115 billion in assets under management (AUM), which is greater than the AUM of all the other robo-advisors combined.[1] PAS is also a hybrid robo-advisor. It features a highly automated investment and planning process and gives investors access to human advisors as part of the service.

We study the evolution of investment advice for a sample of more than 55,000 individuals who were previously self-directed investors and signed up for robo-advice during the 2015–2017 period. Our sample comprises investors with considerable portfolio wealth (median portfolio wealth is $407,652) and a willingness to take equity risk (median equity share 56%). Our primary focus is on the portfolio effects of advice—including portfolio composition and risk-adjusted returns—and the time benefits investors derive from delegating the management of their investment portfolio (see Kim et al., 2016). However, 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 planning, retirement income planning, or benefits plan optimization. Investors may also derive other benefits from advice, such as 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 based on their financial objectives, risk tolerance, investment horizons, and demographic characteristics. They are then offered a comprehensive financial plan, 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 advisor who explains the plan, and investors are officially enrolled in the service only after accepting the proposed plan and agreeing to proceed with the engagement. From that moment, trading occurs automatically on behalf of the investor to reach the desired portfolio allocation. The algorithm revisits investor positions quarterly, and trades are placed if portfolio weights deviate substantially from target weights.[3]

We establish five facts about robo-advice. *Fact 1*: adopting robo-advice materially reshapes self-directed investor portfolios, increasing indexing and reducing home bias, number of assets held, and fees. *Fact 2*: robo-advising, while reducing mean expected returns slightly, also lowers total volatility and idiosyncratic risk even more so, resulting in meaningful improvements in log-Sharpe ratios and, thus, welfare for CARA utility investors. *Fact 3*: the investors who benefit the most from robo-advice are the ones who—before adopting advice—did not have high exposure to equities, did not diversify internationally, and did not invest widely in indexed mutual funds. *Fact 4*: robo-advised investors spend less time monitoring their portfolios, saving on average six hours of time, valued at roughly $450 per year. *Fact 5*: the investors who benefit most from the robo-advice service are more likely to adopt it and less likely to quit it. We detail our findings regarding each fact below.

Starting from *Fact 1*, we document that robo-advice operates significant changes in investors' portfolios. The percentage of wealth in indexed mutual funds almost doubles: it increases from 46% to 81%. Investors' international diversification increases threefold: the percentage of wealth in international mutual funds increases from 11% to 31%. Moving investors into indexed mutual funds translates into lower fees: average expense ratios are more than halved—from 23 to 10 basis points. Robo-advice also changes investors' risky shares. It increases investors' bond holdings from 25% to 39% and decreases investors' cash and money market mutual fund holdings from 19% to 2%. Equity holdings increase, on average, from 56% to 59%.

We also document a substantial reduction in the cross-sectional variation—i.e., homogenization—in portfolio characteristics post-robo-advice adoption. For example, the percentage of equities in investor portfolios ranges from 25% to 100% at the 10th and 90th percentiles before robo-advice adoption. The corresponding percentages post-adoption are 40% and 85%. The effects across other portfolio characteristics are similar.

Turning to *Fact 2*, i.e., investment performance, our main results analyze investors' risk-return trade-off before and after adopting robo-advice. We follow Calvet et al. (2007) and estimate the Sharpe ratio of advised and non-advised investors. In our baseline specifications that use the total return on the MSCI World Index as a benchmark, investors' Sharpe ratios increase by 16.1% after adopting robo-advice. This is mainly due to a reduction in investors' portfolio total risk by 15.8% and idiosyncratic risk by 27.6%. Expected returns, on the other hand, are slightly lower post-robo-advice. 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 performance improvement is economically and statistically significant, starting from the first month after adopting the robo-advisor and persists over time.

One of the limitations of our data is that it allows us to track the effects of robo-advice only on the portion of investors' wealth invested with the asset manager. When we focus on those investors we estimate to allocate the entirety of their financial wealth with the robo-advisor, we find that our results are virtually identical to the ones we estimate on the full sample of robo-advised investors.

The effects of robo-advice computed across all investors hide considerable cross-sectional heterogeneity. To understand which customers are more likely to benefit from robo-advising (*Fact 3*), 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 investment performance pre- and post-advice. We consider 14 covariates that capture a variety of demographic, portfolio, and trading characteristics—all measured before self-directed investors adopt robo-advice. The most important covariates in explaining the cross-sectional variation in the improvement in portfolio performance post-adoption of advice are investors' equity share, cash share, percentage of wealth in mutual funds, and exposure to international mutual funds. This suggests that the investors who benefit the most from robo-advice are the ones who—before adopting advice—did not have high exposure to equities, did not diversify their portfolios internationally, and did not invest widely in indexed mutual funds.

In all cases, however, the relation between investor characteristics and the benefits from robo-advice is nonlinear and, in some cases, nonmonotonic—suggesting that standard statistical methods may lead to inaccurate insights in this context. We perform an out-of-sample cross-validation exercise to show that these nonlinearities are not the result of over-fitting. BRTs do not overfit the training sample and provide superior in- and out-of-sample performance compared to linear models that use the same covariates. 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.

To establish *Fact 4*, we study investor attention before and after signing up for advice. We show that advised investors decrease the

[1] Other large US robo-advisors are Schwab Intelligent Portfolios ($37 billion AUM), Betterment ($16 billion AUM), Wealthfront ($11 billion AUM), and Personal Capital ($8.5 billion AUM).

[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.

[3] In principle, investors can reach out to their advisors to modify their portfolio positions. Unfortunately, we do not observe such occurrences, but anecdotally, they do not happen often.

effort they need to exert to manage their investment portfolios. This reduction in attention is not related to an overall decrease in the investors' awareness of their financial condition because investors tend to log in more often to quickly acquire information regarding their portfolio wealth whenever they need to. Still, the overall time spent making investment decisions decreases after adopting advice. The time savings add up to 6 hours per investor, which we value at roughly $450 per year. Combining these findings with the diversification effects documented in Facts 1 through 3 suggests that investors not only save time by not trading, but they plausibly avoid investment mistakes that would lead to inefficient investing.

Finally, *Fact 5* relates to 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 equity exposure and international diversification, pay high management fees, and have high portfolio volatility as self-directed investors—the most likely to sign up for advice and the least likely to quit the service. The results also suggest that the involvement of human advisors may play a role in increasing the probability of sign-up and decreasing attrition.

Note that, in our setting, the decision to sign up for robo-advice is endogenous, so our results should be interpreted as descriptive rather than causal. At the same time, the changes in portfolio allocations operated by the robo-advisor are mechanical. They are based on a proprietary algorithm that is not known to investors before signing up, so the explanation that appears most plausible is that those investors who sign up for robo-advice let the robo-advisor choose their portfolio allocations on their behalf. Hence, above and beyond the decision to sign up, endogenous investors' characteristics and investor-specific shocks have little role in determining the post-sign-up portfolio composition and performance.

Note also that our main contribution is to establish the qualitative effects of robo-advising on investors' welfare. Establishing the quantitative welfare implications would require formulating and estimating a more complex portfolio choice model and including in our computations additional features such as labor income and the wealth individuals hold outside our asset manager.

1. Related literature

Our work contributes to multiple strands of the finance and economics literature. First, we contribute to the nascent literature in robo-advising. D'Acunto et al. (2019a) are the first to analyze the effects of robo-advising on individual investors' portfolios. They find both promises and pitfalls in that not all customers gain from adopting robo-advising. Our paper differs from D'Acunto et al. (2019a) in many respects. First, the robo-advisor in D'Acunto et al. (2019a) is a portfolio optimizer for stock portfolios. The robo-advisor analyzed here is of a more modern conception. For example, it uses indexed mutual funds rather than individual stocks. A second major difference is that the robo-advisor analyzed here automatically trades for the investor—it is a robo-manager—while the one in D'Acunto et al. (2019a) only recommends portfolio changes to the investor whenever the investor uses it. It is the investor's responsibility to place every trade. Reher and Sun (2019) study the effect of an automated financial management service. They find that automated portfolios are more diversified than self-managed ones and that reducing the minimum balance required to access the service increases customer fund inflows. Our results complement the ones in Reher and Sun (2019) as we assess the risk-adjusted investment performance and the type of portfolios and instruments held pre- and post-adoption of robo-advising. In addition, we analyze the full cross-section of investors to measure which ones benefit the most, the effects of robo-advising on investor attention, and the determinants of robo-advising adoption and attrition. For studies related to the role of trust in robo-advice and the impact of robo-advising on financial inclusion and employee savings plans, see Rossi and Utkus (2019), Reher and Sokolinski (2023) and Bianchi and Briere (2020). Finally, for a comprehensive review of the robo-advising literature, see D'Acunto and Rossi (2021).

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 cannot make him drink*". Our results stress that the automatic implementation of advice is crucial for the efficacy of any form of financial advice, whether human- or robo-generated.

Second, and more broadly, we contribute to the household finance literature. Campbell (2006) argues financial markets benefit households only to the extent that they participate in financial markets and hold instruments that provide them with well-diversified investment portfolios. As shown in Badarinza et al. (2016), although households with higher socioeconomic 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 et al., 2019b). Relatedly, Guiso et al. (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 them. Our results show that 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 et al., 2015). However, for many retail investors traditional financial advisors are too costly. In addition, using data from the Canadian advice market, Linnainmaa et al. (2018) show that the increased risk-taking on the part of the clients does not compensate for the higher costs associated with employing a financial advisor. Moreover, advisors often adopt a one-size-fits-all approach and might be prone to behavioral biases or display cognitive limitations (Linnainmaa et al., 2021). Fintech robo-advising gives clients access to financial advice at a 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 human advisors.

Our study is also relevant to the broader literature on technology adoption. Romer (1990) and Aghion and Howitt (1992) maintain that the adoption of new technologies is a crucial determinant of economic growth. Comin and Mestieri (2014) argue there is a lack 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 setting and data sources used in the study. We then present the demographic and portfolio characteristics of robo-advised investors before they sign up for the 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. Setting and data

The study uses anonymized proprietary data from Vanguard's Personal Advisor Services (PAS). Because we analyze a specific robo-advisor, we cannot argue that the effects we uncover extend to the whole robo-advising industry, even though anecdotal evidence suggests that most mainstream robo-advisors are similar from a design perspective. An advantage of our setting is that the robo-advisor we analyze is the largest in the world, and it has more AUM than the rest of the US robo-advisors combined, as highlighted in the introduction. In this sense, our results are highly representative of the robo-advising industry.

Our data contains information on trades, positions, demographic characteristics, and investor–advisor mapping information for previously self-directed investors who have interacted with the robo-advisor. The main results are computed using the sample of previously self-directed investors who signed up for advice between January 2015 through December 2017.

The trade data includes all trades from January 2015 through December 2017. The position data contains monthly holdings observations for the same investors. The demographic characteristic data contains information on investors' age and gender and detailed information regarding the dates investors initiated, enrolled, implemented, and quit the advice service. The investor–advisor mapping data contains information on the interactions between investors and human advisors, i.e., meetings, phone calls, etc.

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 the demographic and portfolio characteristics of the investors who sign up for advice, computed the month before the investors sign up for the service. We restrict the analysis to those investors who remain with the service for at least six months to facilitate comparisons across tables.[4] 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 10th to the 90th percentile. Panel A focuses on the demographic characteristics. The average investor is 64 years old (the median is 65), and 60% of the users are males. Tenure as self-directed investors varies widely. It ranges from two years at the 10th percentile to 26 years at the 90th percentile. For comparison, the average investor age is 51 in Gargano and Rossi (2018) and Barber and Odean (2001). The percentage of women, which equals 40%, is larger compared to both (Gargano and Rossi, 2018), 27%, and Barber and Odean (2001), 21%. At approximately 15 years, the average investor tenure is also longer compared to other brokerage account datasets in the literature. Average investor tenure in Gargano and Rossi (2018) is less than 9 years.

Panel B of Table 1 reports results for portfolio allocation. Investors' wealth is substantial. It averages $723,010 and is heavily skewed to the right. The median invested wealth is $407,652. The number of assets per investor is 10.79, and the median is 7. It may appear that these investors are substantially under-diversified, but this is really not the case because, as we show below, the majority of these investors are very heavily invested in mutual funds.

The average investor has 56% of their portfolio invested in equities, followed by 25% in bonds, and 19% 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 a small fraction of the investors invested only in bonds and/or cash. Stocks and bonds are not held directly but mainly through mutual funds. In fact, 73% of the wealth is invested in mutual funds, followed by cash at 17%. Only 4% of investors' wealth is held in individual stocks and 4% 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, 46% of mutual fund holdings are in indexed mutual funds, and 11% of mutual fund holdings are in funds that invest internationally. Only a negligible number of investors 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 16 basis points, but the 90th percentile of investors spend as much as 35 basis points a year on management fees. The expense ratio results are similar. The average is 0.23, the median is 0.15, 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, which averages 0.35.

Finally, investors place on average 5 transactions per month, for an average of $145K, but the median is much smaller at only $2.4K. As we show below, these quantities do not represent the steady-state level of investor activity because investors make more transactions in the months immediately preceding signing up for advice. They generally transact in an effort to consolidate their accounts before enrolling in the service.

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 sixth month post-adoption of advice. The demographic characteristics such as age, tenure, and proportion of males—in Panel A of Table 2—are reported for completeness. Still, they are identical to those reported in Table 1 because we condition on the same sample of investors. Panel B reports the portfolio allocation results. At $819,519, average wealth is higher than in Table 1. This is the result of stock market appreciation and investors' contribution to their portfolios. The number of assets decreases slightly from 10.79 to 8.62. The percentile distribution shows that advice shrinks the number of stocks held in the tails of the distribution. The 90th percentile of the number of assets held in each account drops from 23 in Table 1 to 15 in Table 2.

Continuing with the results in Panel B, we observe large changes in portfolio allocation, particularly in the allocation to bonds and money market mutual funds (cash). The percentage of bonds increases by 14 percentage points to 39%, while the allocation to cash decreases by 17 percentage points to only 2%. Finally, the equity share increases by three percentage points to 59%. The following four lines in Panel B of Table 2 focus on the investment vehicles used. Almost all of investors' wealth—94% of it—is invested in mutual funds, with virtually no share of wealth in money market mutual funds (2%), ETFs (2%), 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 81% after signing up for advice. We find a similar effect for investors' exposure to international markets that increases from 11% to 31%. Interestingly, we do not detect much of an effect in terms of emerging markets exposure that, if anything, declines after advice. As shown in Online Appendix A, this is because international mutual funds (VTIAX, for example) have an emerging market exposure that the Lipper classification does not capture well.

[4] We would like to thank an anonymous referee for this suggestion.

[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.

Table 1
Demographic and portfolio characteristics of Robo-advised investors 1 month before Sign-up.

	N	mean	sd	p10	p25	p50	p75	p90
Panel A. Demographic Characteristics								
Age	55,202	63.79	12.66	46.00	57.00	65.00	72.00	79.00
Male	55,202	0.60	0.49	0.00	0.00	1.00	1.00	1.00
Tenure	55,202	15.04	9.15	2.42	6.67	15.67	21.17	26.50
Panel B. Portfolio Allocation								
	N	mean	sd	p10	p25	p50	p75	p90
Wealth ($)	55,202	723,010	892,523	84,978	174,767	407,652	889,005	1,703,176
NumAssets	55,202	10.79	10.28	2.00	4.00	7.00	13.00	23.00
PctEquityShare	55,202	0.56	0.27	0.15	0.38	0.58	0.76	0.92
PctBondShare	55,202	0.25	0.22	0.00	0.05	0.21	0.40	0.56
PctCashShare	55,202	0.19	0.28	0.00	0.00	0.04	0.27	0.68
PctMutualFunds	55,202	0.73	0.32	0.17	0.52	0.90	1.00	1.00
PctCash	55,202	0.17	0.26	0.00	0.00	0.03	0.23	0.61
PctStocks	55,202	0.04	0.11	0.00	0.00	0.00	0.01	0.16
PctETF	55,202	0.04	0.12	0.00	0.00	0.00	0.00	0.13
PctIndex	55,202	0.46	0.34	0.00	0.14	0.45	0.77	0.97
PctInternational	55,202	0.11	0.13	0.00	0.00	0.06	0.18	0.29
PctEmerging	55,202	0.01	0.02	0.00	0.00	0.00	0.00	0.02
Panel C. Transactions and Fees								
	N	mean	sd	p10	p25	p50	p75	p90
MgtFee	55,202	0.16	0.14	0.04	0.07	0.12	0.20	0.35
ExpRatio(×100)	55,202	0.23	0.21	0.04	0.10	0.15	0.25	0.51
TurnRatio	55,202	0.35	0.28	0.09	0.16	0.28	0.43	0.67
Transaction	55,202	5.21	9.25	0.00	0.00	1.00	5.00	16.00
Volume ($)	55,202	144,966	328,307	0.00	0.00	2,371	100,000	486,666

This table reports the demographic characteristics and portfolio allocation behavior the month before signing up for advice for investors that stay with robo-advice for at least six months after signing up. 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; $Tenure$, 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 10th, 25th, 50th, 75th, and 90th.

Panel C of Table 2 shows that robo-advice moves investors into passive mutual funds with lower fees and turnover ratios. Management fees are halved, from 16 to 8 basis points, while the expense ratio is reduced by more than 50% as it drops from 23 to 10 basis points. The turnover ratio instead drops by approximately 23%, from 0.35 to 0.27.

The results in Tables 1 and 2 present snapshots one month before and six months after robo-advising adoption. In Figs. 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 the median values, while the red dashed lines represent the distribution's 10th and 90th percentiles. Time "0" denotes the month before investors sign up for advice.

Subfigures (a), (b), (c), and (d) of Fig. 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 features. First, it takes one to two months for the service to converge to the new portfolio allocations post-adoption. Second, investors tend to modify their investment portfolio in the months leading to signing up for advice. This is most evident in the results related to the percentage wealth investors have in cash and mutual funds (Subfigures (b) and (d)). In particular, we observe an increase in cash holdings that is not very large at the median but rather pronounced at the 90th percentile, and we observe a decrease in mutual fund holdings that is, once again, not very pronounced at the median but rather large at the 10th percentile. This has—potentially—implications when evaluating investors' performance pre- and post-advice, as we discuss in Section 3.1. Third, the results show a substantial reduction in the cross-sectional variation—i.e., homogenization—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)), 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.

Fig. 2 presents results for indexation, international diversification, expense ratio, and trading volume. In all cases, the changes take place over one or two months, and we observe a substantial reduction in the cross-sectional dispersion across investors when it comes to indexation, percentage of wealth in international funds, and expense ratios, reported in Subfigures (a) through (c). The trading volume results (Subfigure (d)) are unique as they display marked non-monotonicities. Trading volume spikes for approximately one to two months after enrollment into the service as the advice service changes investors' positions to the new target weights.

Finally, when we test for differences in the average value for each portfolio characteristic reported in Figs. 1 and 2 before and after adopting robo-advice, we find that the differences are statistically significant with p-values smaller than 0.1% in all cases.

Online Appendix A.1 presents a detailed analysis of the actual tickers held by advised and non-advised investors. Advised investors have

Table 2
Demographic and portfolio characteristics of advised investors 6 months after sign-up.

	N	mean	sd	p10	p25	p50	p75	p90
Panel A. Demographic Characteristics								
Age	55,202	63.79	12.66	46.00	57.00	65.00	72.00	79.00
Male	55,202	0.60	0.49	0.00	0.00	1.00	1.00	1.00
Tenure	55,202	15.04	9.15	2.42	6.67	15.67	21.17	26.50
Panel B. Portfolio Allocation	N	mean	sd	p10	p25	p50	p75	p90
Wealth	55,202	819,519	952,378	113,089	219,719	494,323	1,020,442	1,886,003
NumAssets	55,202	8.62	5.25	4.00	5.00	7.00	10.00	15.00
PctEquityShare	55,202	0.59	0.18	0.39	0.49	0.58	0.70	0.85
PctBondShare	55,202	0.39	0.18	0.13	0.28	0.40	0.49	0.59
PctCashShare	55,202	0.02	0.05	0.00	0.00	0.00	0.01	0.06
PctMutualFunds	55,202	0.94	0.11	0.80	0.94	1.00	1.00	1.00
PctCash	55,202	0.02	0.05	0.00	0.00	0.00	0.01	0.06
PctStocks	55,202	0.02	0.05	0.00	0.00	0.00	0.00	0.05
PctETF	55,202	0.01	0.04	0.00	0.00	0.00	0.00	0.02
PctIndex	55,202	0.81	0.18	0.54	0.72	0.85	1.00	1.00
PctInternational	55,202	0.31	0.09	0.18	0.27	0.34	0.36	0.39
PctEmerging	55,202	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	55,202	0.08	0.03	0.06	0.06	0.07	0.08	0.11
ExpRatio ($\times$100)	55,202	0.10	0.03	0.07	0.08	0.09	0.10	0.14
TurnRatio	55,202	0.27	0.12	0.09	0.19	0.28	0.34	0.40
Transaction	55,202	2.65	4.13	0.00	0.00	1.00	3.00	8.00
Volume ($)	55,202	16,836	55,373	0.00	0.00	259	3,690	30,214

This table reports the demographic characteristics and portfolio allocation behavior of investors 6 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; $Tenure$, 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 twelfth month after 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 10th, 25th, 50th, 75th, and 90th.

portfolio allocations that are more homogeneous—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).

We summarize this section's findings in *Fact 1*, reported below:

Fact 1*: Adopting robo-advice materially reshapes self-directed investor portfolios, increasing indexing and reducing home bias, number of assets held, and fees.*

3. Performance before and after robo-advice

The portfolio allocation results reported so far suggest that the advice service may improve investors' performance by placing account holders in diversified US and international low-fee indexed mutual funds. It also reduces investors' cash holdings. In this section, we first provide a comprehensive analysis of how robo-advice adoption relates to investment performance (Section 3.1). We then show the robustness of our results when we focus on individuals estimated to hold most of their investable assets with the robo-advisor (Section 3.2).

3.1. Robo-advising and investment performance

In this section, we follow Calvet et al. (2007) and Reher and Sokolinski (2023) and estimate the Sharpe ratio of advised and non-advised investors by first obtaining returns data on the universe of all assets from January 1990 until December 2017. We then estimate a simple CAPM model for all individual securities as follows:

$$R_{k,m} = \beta_k F_m + \epsilon_{k,m} \tag{1}$$

where $R_{k,m}$ is the return of security k in month m in excess of the risk-free rate, β_k is the loading on F_m—the total return on the MSCI World Index, which has an average annual excess return of 5.3% over our sample period. We work with a simple CAPM specification, given the growing evidence that investors' CAPM alphas are the best predictor of mutual flows—see Berk and Van Binsbergen (2016) and Barber et al. (2016).[6]

In the second step, we compute portfolios' overall risk as well as expected returns and idiosyncratic risk and Sharpe ratios on the basis of this linear factor model while focusing on investors' portfolio

[6] Multi-factor specification delivers qualitatively similar results.

Table 3
Performance comparison between self-directed and robo-advised portfolios.

	N	Mean	SD	p5	p25	p50	p75	p95
Panel A. All Investors								
Expected Returns								
Self-directed	38,049	3.22%	1.36%	0.19%	2.46%	3.36%	4.24%	5.07%
Robo-advised	38,874	3.03%	1.04%	1.28%	2.41%	2.99%	3.66%	4.81%
Volatility								
Self-directed	38,049	9.29%	3.89%	2.27%	6.93%	9.47%	11.99%	14.50%
Robo-advised	38,874	7.82%	2.55%	3.85%	6.19%	7.57%	9.32%	12.41%
Idiosyncratic Volatility								
Self-directed	38,049	2.93%	2.16%	0.90%	1.95%	2.42%	3.43%	5.99%
Robo-advised	38,874	2.12%	1.10%	1.50%	1.66%	1.87%	2.24%	3.40%
Sharpe Ratio								
Self-directed	38,049	0.329	0.094	0.078	0.334	0.349	0.365	0.404
Robo-advised	38,874	0.382	0.057	0.314	0.370	0.398	0.411	0.424
Panel B. Investors with the Majority of their Assets in the Robo-advisor								
Expected Returns								
Self-directed	7,901	3.22%	1.32%	0.41%	2.44%	3.35%	4.20%	5.05%
Robo-advised	8,006	3.00%	0.96%	1.45%	2.40%	2.94%	3.56%	4.71%
Volatility								
Self-directed	7,901	9.27%	3.79%	2.78%	6.87%	9.44%	11.89%	14.41%
Robo-advised	8,006	7.74%	2.39%	4.10%	6.19%	7.53%	9.13%	12.02%
Idiosyncratic Volatility								
Self-directed	7,901	2.93%	2.15%	1.09%	1.96%	2.41%	3.39%	6.10%
Robo-advised	8,006	2.11%	1.04%	1.50%	1.66%	1.87%	2.22%	3.27%
Sharpe Ratio								
Self-directed	7,901	0.334	0.087	0.172	0.335	0.350	0.366	0.404
Robo-advised	8,006	0.385	0.045	0.325	0.369	0.397	0.411	0.424

This table reports portfolio expected returns, volatility, idiosyncratic volatility, and Sharpe Ratios for robo-advised and self-directed investors, computed following (Calvet et al., 2007) as described in Section 3.1 and using the MSCI World Index as a benchmark and focusing on investors' portfolio holdings at the sixth month before and after adopting advice. For each performance metric, we report the cross-sectional average, standard deviation, 5th, 25th, 50th, 75th, and 95th percentiles. Panel A reports the results for all investors. As described in detail in Section 3.2, Panel B repeats the computations for the subset of investors who are estimated to have the majority of their wealth invested in the robo-advisor.

holdings six months before and after adopting advice. We focus on the sixth month before and after robo-advice because, as we show in Section 2.3, investors tend to increase their cash holdings in the months preceding the adoption of the service, and this may negatively affect the performance of their investment portfolio. Our choice is, therefore, conservative. We focus on the investors for whom we can match 90% of their holdings to individual asset returns, but the results are not sensitive to this choice.

We report the results in Panel A of Table 3 that reports the cross-sectional average, standard deviation, and the 5th, 25th, 50th, 75th, and 95th percentiles of the distribution of each performance metric, computed separately for self-directed and advised portfolios. The expected returns of advised portfolios are similar—albeit slightly smaller—to those of self-directed portfolios (3.22% versus 3.03%), while the risk of advised portfolios is $1 - 7.82\%/9.29\% = 15.8\%$ lower than the self-directed ones. Part of the reduction in risk is due to a reduction in idiosyncratic risk that is 27.6% lower for advised investors (2.12%) compared to self-directed investors (2.93%). As a result, the Sharpe Ratios of the advised portfolios are 16.1% larger than the ones of the self-directed portfolios (0.382 versus 0.329). Comparing the Sharpe ratios of self-directed and advised portfolios shows that the latter are superior at every reported percentile and have a much lower cross-sectional dispersion. For example, the Sharpe ratios for self-directed investors range from 0.078 for the 5th percentile to 0.404 for the 95th, compared to a range of 0.314 to 0.424 for advised investors.

The results, reported in Table 3, confound variation across and within individuals, so we re-estimate our results using individual-level changes in log Sharpe ratios. We do so because changes in log Sharpe ratios difference out the effect of market expected returns and, for investors with CARA-normal utility, the change in log Sharpe ratio relative to a self-managed portfolio maps directly to the welfare gain from robo-advice.[7] We report our results in Subfigure (a) of Fig. 3. Most investors improve their portfolio allocation after adopting robo-advice, as the density of the differences in log Sharpe ratios is well above zero. The differences are economically large and statically significant as, on average, investors experience a 10.9% increase in their log Sharpe ratios (blue vertical line) when robo-advised compared to when they self-manage their portfolios, with an associated t-statistic of 77. We find similar results when we repeat the analysis for risk and idiosyncratic risk. For example, for the within-investor log difference in idiosyncratic risk after adopting robo-advice, compared to before,

[7] We thank an anonymous referee for making this point.

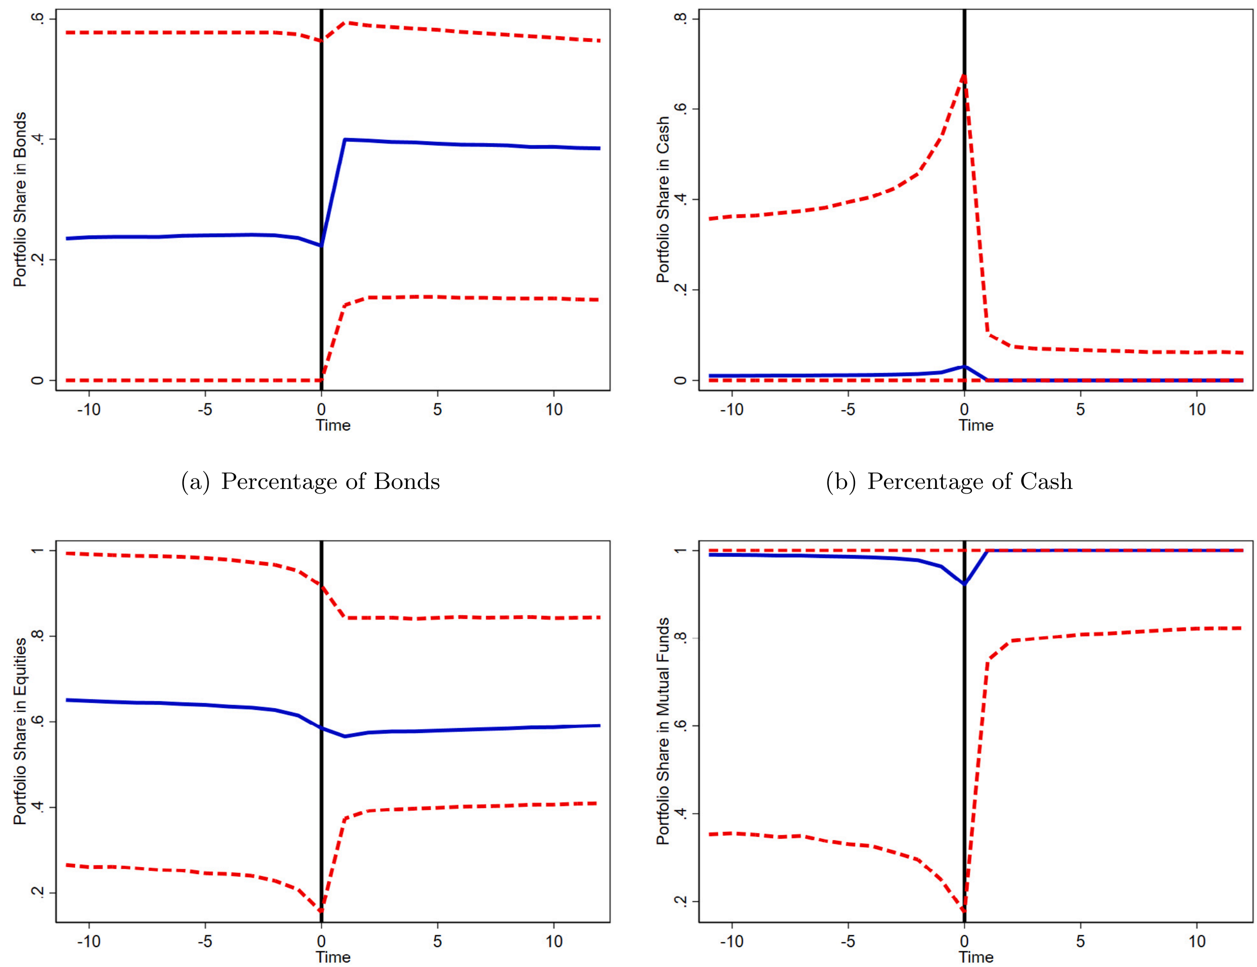

(a) Percentage of Bonds

(b) Percentage of Cash

(c) Percentage of Equities

(d) Percentage in Mutual Funds

Fig. 1. Portfolio allocation before and after advice. 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. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

we find a reduction of 17.5%, with a very large (negative) t-statistic equal to −48.6.

Using different asset pricing models delivers qualitatively very similar results. For example, if we use the market value-weighted return on the NYSE, AMEX, and Nasdaq downloaded from CRSP, we obtain higher expected Sharpe ratios as a result of the higher expected returns for both self-directed and robo-advised portfolios, but the significant improvement associated with adopting robo-advice remains.[8] Self-directed investors have Sharpe ratios equal to 0.499, while robo-advised investors have Sharpe ratios equal to 0.561, statistically different from each other at the 1% level. Economically, this implies a risk-return trade-off improvement of 0.561/0.499=12.4%. Finally, in Subfigure (b) of Fig. 3, we report the density of the differences in log-Sharpe ratios for this alternative asset pricing model. The distribution is overwhelmingly positive, and the differences are economically large and statically significant as, on average, investors experience a 7.4% increase in their Sharpe ratios when robo-advised compared to when they self-manage their portfolios, with an associated t-statistic equal to 52.7.

Note that all the results reported so far focus on investors' portfolio holdings six months before and after adopting advice. We obtain similar results if we repeat the analysis at different horizons. We show this in Subfigures (c) and (d) of Fig. 3, where we estimate dynamic specifications using all the holdings in the 12-month window around robo-advice adoption:

$$Sharpe_{i,t} = \alpha_i + \beta_t + \sum_{j=-3}^{5} \gamma_j \ ROBO_{i,j,t} + \epsilon_{i,t}, \tag{2}$$

where $Sharpe_{i,t}$ is the Sharpe Ratio computed following Calvet et al. (2007), α_i denote investors' fixed effects, β_t are time-effects, and the dummy variable $ROBO_{i,j,t}$ is equal to zero for every month, except for the jth month before and after adoption. The 0th month is 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 additional advantage of working with these dynamic regressions is the ability to control for both individual and time effects. We report coefficients and 95% confidence intervals based on double-clustered standard errors using the international benchmark in Subfigure (c) and the US benchmark in Subfigure (d).

[8] The value-weighted index on the NYSE, AMEX, and Nasdaq has an average annual excess return of 7.9% over our sample period.

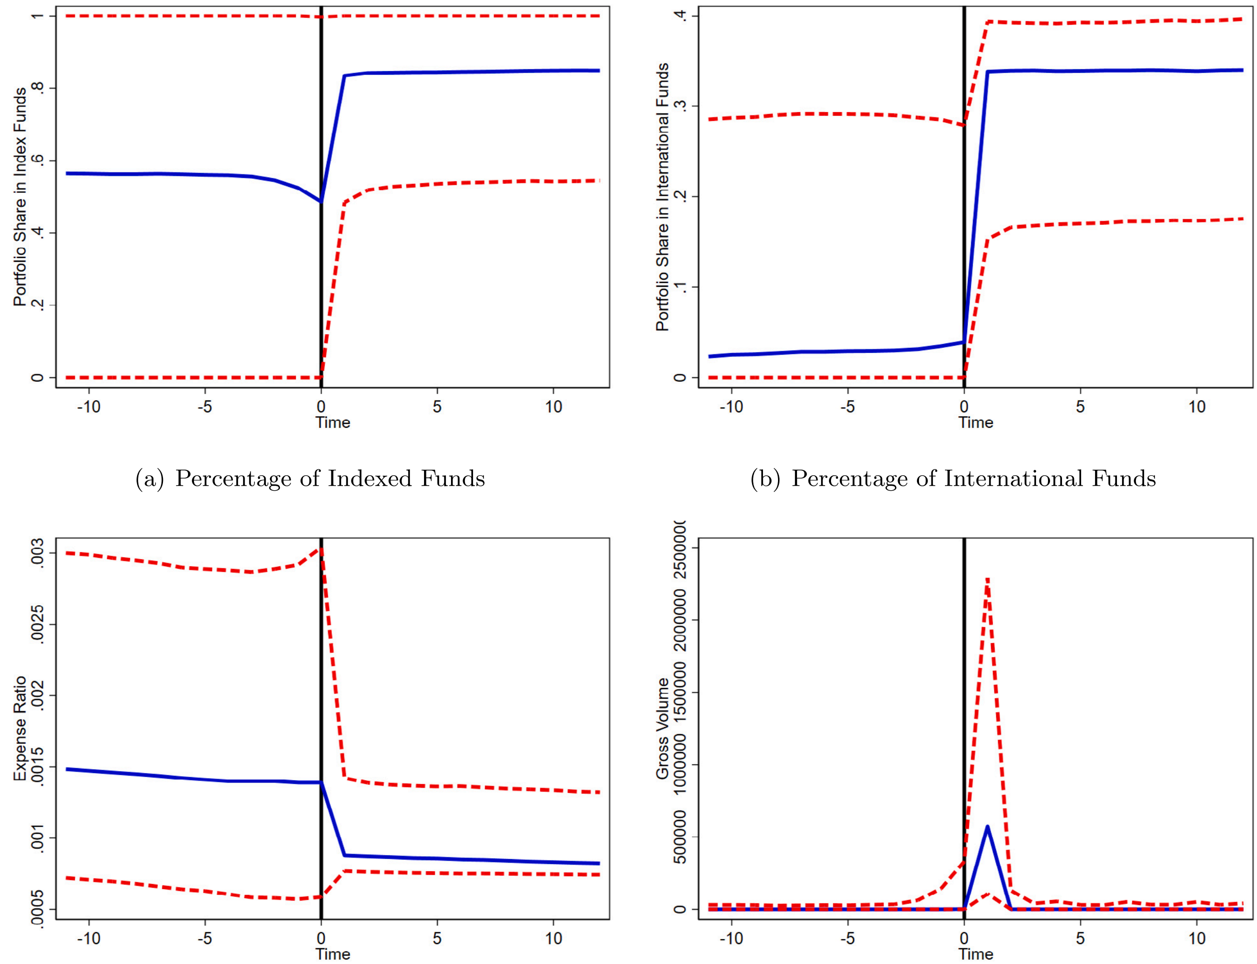

(a) Percentage of Indexed Funds

(b) Percentage of International Funds

(c) Expense Ratios

(d) Trading Volume

Fig. 2. Indexation, international diversification and fees before and after advice. 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. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

In both cases, we see significant and persistent improvement in risk-adjusted performance from robo-advice adoption. The plot also shows that, in the three months preceding the adoption of robo-advice, investors' performance suffers slightly from the fact that they increase their cash holdings. This can be seen from the fact that the coefficient estimates for periods −1, −2, and −3 are negative and marginally statistically significant at the 95% levels.

3.2. Accounting for total wealth

One of the main advantages of obtaining data from individual asset managers or fintech apps is that we can observe individuals' trading at a high frequency. On the other hand, country-level administrative data, like the Swedish or Danish datasets used in many household portfolio studies (Campbell, 2006), only track individual behavior at a yearly frequency. The drawback of fintech data compared to country-level data, on the other hand, is that it is not possible to keep track of the entirety of investors' wealth. As a result, some of the investment performance improvements we document may be undone by investors' behavior in other accounts held at different asset managers.

In this section, we overcome this limitation using an additional variable that estimates investors' total investable wealth.[9] When computing the ratio of wealth invested in the robo-advisor and dividing it by the overall investors' wealth, we find that the average value of the ratio is 90% and the median is instead 55%, indicating that the measure is somewhat right-skewed. At the 90th percentile, investors' assets in the

[9] We use WealthComplete data from Equifax. The data provides a Total Assets measure, expressed as a continuous measure of estimated dollars per household of all deposits plus investments capped at a maximum value of $25 Million+. These estimates include personal financial investments held in taxable, IRA, and Keogh accounts, including deposits, investments, and annuities, and exclude assets held in 401K, 403B, profit sharing, IRA-SEP, stock purchase/ESOP, money purchase plans, business accounts, life insurance, or home value. The wealth information is available at the ZIP+4 digits level, allowing us to match it with our users' data with high precision.

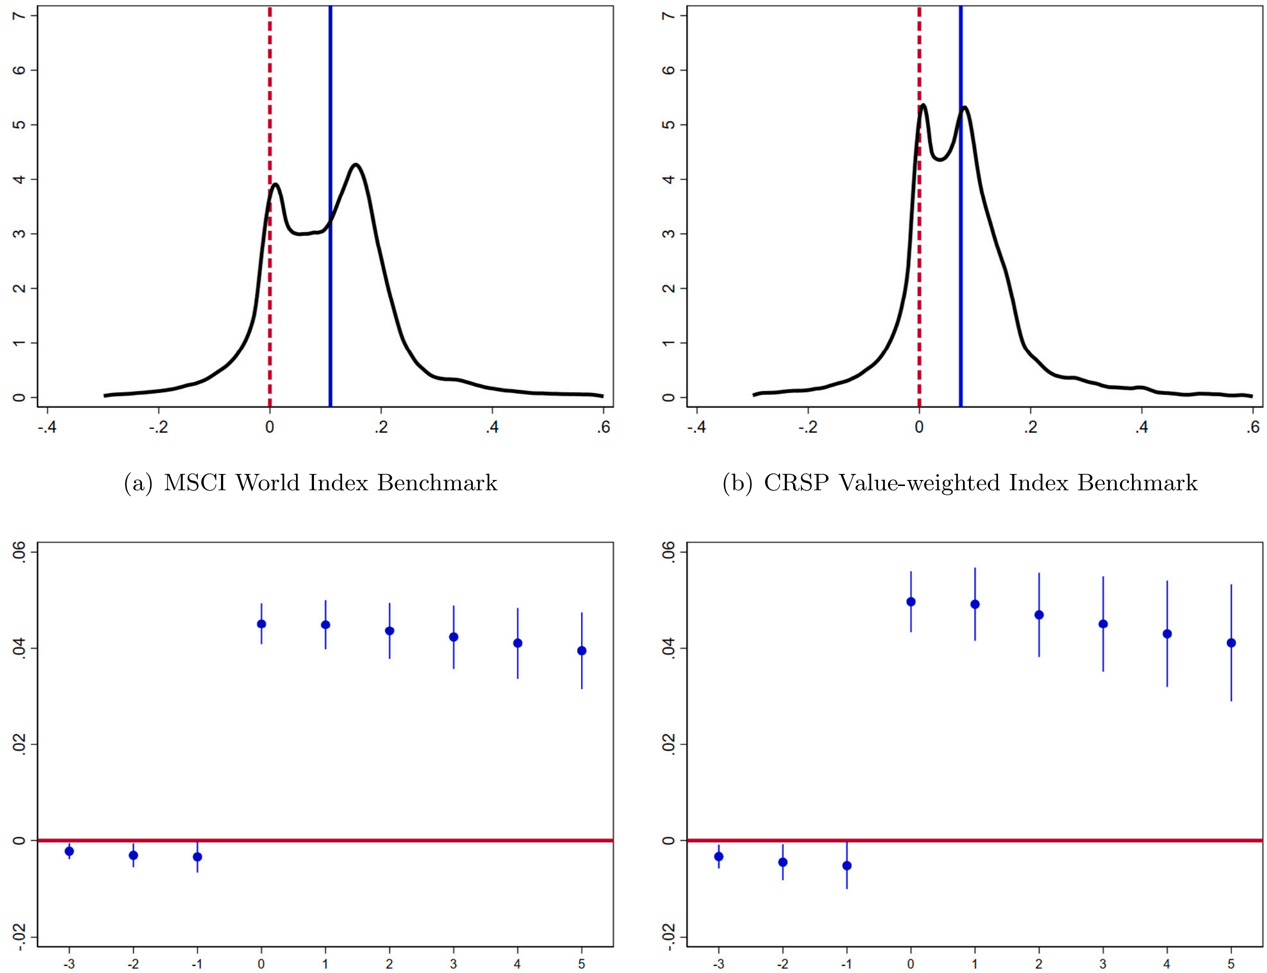

(a) MSCI World Index Benchmark

(b) CRSP Value-weighted Index Benchmark

(c) MSCI World Index Benchmark

(d) CRSP Value-weighted Index Benchmark

Fig. 3. Changes in investors' sharpe ratios before and after adopting robo-advice. Subfigure (a) reports the density of the within-investor differences in log-Sharpe ratios after compared to before signing up for robo-advice. Investors' Sharpe ratios are computed following Calvet et al. (2007) as described in Section 3.1 and using the MSCI world index as a benchmark. The vertical blue line represents the average of the distribution. Subfigure (b) repeats the exercise using the NYSE, AMEX, and Nasdaq CRSP value-weighted index. Subfigure (c) reports estimates of the following dynamic regression specifications:

$$Sharpe_{i,t} = \alpha_i + \beta_t + \sum_{j=-3}^{5} \gamma_j \ ROBO_{i,j,t} + \epsilon_{i,t},$$

where $Sharpe_{i,t}$ is the Sharpe Ratio computed following Calvet et al. (2007) as described in Section 3.1 and using the MSCI world index as a benchmark, α_i denote investors' fixed effects, β_t are time-effects, and the dummy variable $ROBO_{i,j,t}$ is equal to zero for every month, except for the jth 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. The results are computed using all holdings in the six months before and after the adoption of robo-advice. Standard errors are double-clustered by user and time. Subfigure (d) repeats the exercise using the NYSE, AMEX, and Nasdaq CRSP value-weighted index. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

robo-advisor equal 240% of their estimated wealth, suggesting that the measure is imperfect.

To be conservative, we limit our computations to the subset of investors whose assets invested with the robo-advisor are between 75% and 125% of their estimated total assets, as those should be the ones for which the concerns reported above should be minimized. We then repeat the analysis reported in Section 3.1 for this subset of investors and report the results in Panel B of Table 3.

The results for this subset of investors are virtually identical to those computed on our full sample. The expected returns of robo-advised portfolios are close to those of self-directed portfolios (3.22 versus 3.00). In contrast, the risk of robo-advised portfolios is $1 - 7.74\%/9.27\% = 16.5\%$ lower than the self-directed ones. As a result, the Sharpe Ratios of the robo-advised portfolios are 15.3% larger than those of the self-directed portfolios (0.385 versus 0.334). In this sample, however, we are more confident that the wealth invested with the robo-advisor is a substantial fraction, if not the entire fraction, of individuals' investable wealth. We can, therefore, be reasonably assured that the improvements we document may not be undone by investors' behavior in other accounts held at different asset managers.

We summarize this section's findings in *Fact 2*, reported below:

Fact 2: *Robo-advising, while reducing mean expected returns slightly, also lowers total volatility and idiosyncratic risk even more so, resulting in meaningful improvements in log-Sharpe ratios and, thus, welfare for CARA utility investors.*

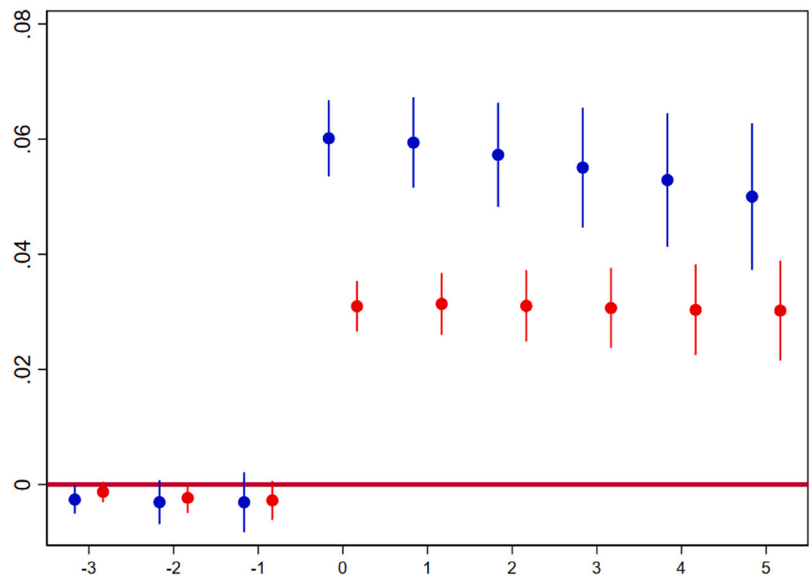

Fig. 4. Changes in investors' sharpe ratios conditioning on their equity holdings before robo-advice. This figure reports estimates of the following dynamic regression specification:

$$Sharpe_{i,t} = \alpha_i + \beta_t + \sum_{j=-3}^{5} \gamma_j \ ROBO_{i,j,t} + \epsilon_{i,t},$$

where $Sharpe_{i,t}$ is the Sharpe Ratio computed following Calvet et al. (2007) as described in Section 3.1 and using the MSCI world index as a benchmark, α_i denote investors' fixed effects, β_t are time-effects, and the dummy variable $ROBO_{i,j,t}$ is equal to zero for every month, except for the jth 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. The results are computed using all holdings in the six months before and after the adoption of robo-advice. Estimates associated with investors with equity holdings below the median before signing up for the robo-advice are reported in blue. The ones associated with investors with equity holdings above the median are reported in red. Standard errors are double-clustered by user and time. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

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 investor. However, we can expect the robo-advisor to have a differential impact on the portfolio allocation and the performance of each investor, depending on investor characteristics at sign-up. To illustrate the point, we report in Fig. 4 the estimates associated with Eq. (2) for investors with equity holdings below the median (in blue) and above the median (in red) before signing up for the robo-advice. The plot shows that the former group experiences a much more significant improvement in its risk-return trade-off than the latter.

The portfolio changes operated by the robo-advisor are primarily 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, and younger individuals a higher one. Investors' lifestyles may also play a role: 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 and the algorithm used by the robo-advisor. It is, therefore, challenging to know what factors ultimately play a role.

A standard way to analyze this problem would be to use linear regression. Still, it is not clear that investors' demographic and portfolio characteristics are linearly related to the changes in investors' portfolios. It is also unclear *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 estimating spurious relations between regressors and regressand. Instead, we use a machine learning method known as Boosted Regression Trees. Boosted Regression Trees not only allow large conditioning information sets but also allow for non-linearities—all without overfitting or falling prey to the so-called curse of dimensionality. Below, we provide an intuitive introduction to Boosted Regression Trees, partial dependence plots, and relative influence measures. We refer the reader to Online Appendix B.1 for a more formal treatment of the topic.

4.1. An intuitive introduction to Boosted Regression Trees

Boosted Regression Trees are a combination of Regression Trees and Boosting. Regression Trees are a machine-learning method that models any continuous dependent variable using recursive binary partitions of the space spanned by the covariates available and modeling the dependent variable as a constant in each of these partitions. In a simple univariate case where the researcher has one dependent variable y, one independent variable x, and wants to fit a tree with only one split, the procedure would select the optimal threshold x_1 such that by modeling the dependent variable as a constant $\bar{y}_1$ to the left of the threshold and a constant $\bar{y}_2$ to the right of the threshold, the empirical error of the model is minimized. If the researcher wanted to feature two splits instead, the procedure would select a second threshold x_2 and model the dependent variable as three distinct constants $\bar{y}_1$, $\bar{y}_2$, and $\bar{y}_3$ in the three regions generated by the splits x_1 and x_2. 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 unclear which variables may be more or less relevant *ex-ante*. Furthermore, it is difficult to know whether there is a linear relation between predictor and predicted variables in our context.

One limitation associated with regression trees is that the approach is sequential, and successive splits are performed on fewer and fewer observations, increasing the risk of fitting idiosyncratic data patterns. Furthermore, the sequential splitting algorithm is not guaranteed to lead to the globally optimal solution. To address these limitations, the machine learning literature frequently pairs Regression Trees with Boosting.

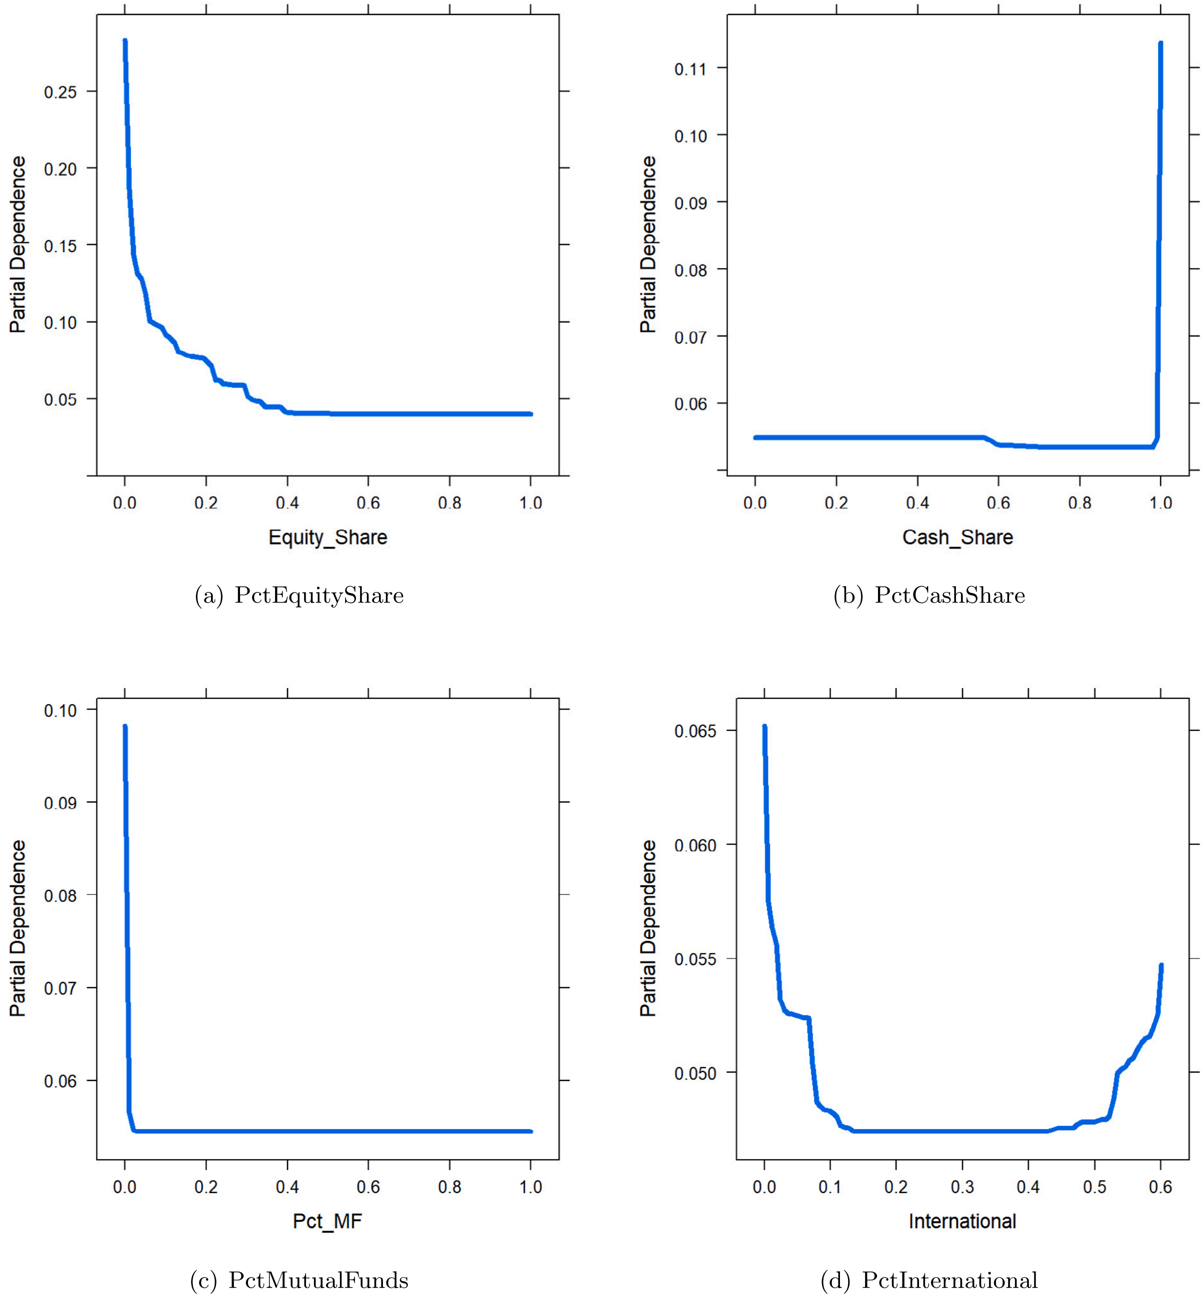

(a) PctEquityShare (b) PctCashShare

(c) PctMutualFunds (d) PctInternational

Fig. 5. Partial dependence plots for investment performance changes post-advice. This figure presents the partial dependence plots for the change in Sharpe Ratios as a function of the 14 regressors described in Section 4.2. In Subfigures (a) through (d), we report partial dependence plots for the 4 predictor variables with the highest relative influence: $PctEquityShare$, the share of wealth in Equities (relative influence of 78.1%); $PctCashShare$, the share of wealth in cash—including money market mutual funds (relative influence of 8.5%); $PctMutualFunds$, the percentage of wealth directly invested in mutual funds (relative influence of 5.8%); and $PctInternational$, the percentage of mutual fund wealth invested in global or international funds—identified using the Lipper mutual fund classification (relative influence of 2.9%). The horizontal axis covers the sample support of each predictor variable, while the vertical axis tracks the change in Sharpe ratios as a function of each individual predictor variable.

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 of observations modeled poorly by the existing collection of trees. By featuring subsampling, model combination, and shrinkage, BRTs routinely rank at the very top of machine learning algorithms when it comes to prediction accuracy, particularly when the relation between regressors and regressand is non-linear, and some of the covariates included in the model are not related to the dependent variable (Gu et al., 2020).

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 such as relative influence measures and partial dependence plots. Relative influence measures provide a ranking of regressors based on how often they are used in the trees and how much they improve the model. Relative influence measures are scaled in such a way that the total value across all regressors sums to 100, and a higher value means that the regressor has a more significant influence on the model's predictions.

Partial dependence plots provide insights into how a particular regressor affects the predicted outcome. They show the relationship between each regressor and the dependent variable, integrating out the effect of all the other regressors.

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.[10] To decompose and visualize how robo-advice changes investors' performance after sign-up, we estimate a BRT model with 10,000 boosting iterations. The dependent variable is the change in the risk-return trade-off (measured by the Sharpe ratio) before and after signing up for advice, using portfolio characteristics measured in the sixth month before and after robo-advising adoption. As conditioning variables, we use a total of 14 regressors divided into three groups. The first group contains demographic characteristics: Age, the investor's age as of December 2017; $Male$, whether the investor is male; $Tenure$, 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; $PctCashShare$, 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; $PctInternational$, the percentage of mutual fund wealth invested in global or international funds—identified using the Lipper mutual fund classification. The third group 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 in a month before signing up for advice; $Volume$, the volume (in US dollars) traded by the investors in a month before signing up for advice.

We report the relative influence measures in the first column of Table 4. Out of the 14 regressors, only 5 covariates have a relative influence higher than 2%. The variable $PctEquityShare$ has the highest relative influence measure, totaling 78.1%. This means that the splits based on $PctEquityShare$ contribute to 78% of the reduction in the empirical error of the model. The second variable is $PctCashShare$ which has a relative influence of 8.5%, followed by $PctMutualFunds$, $PctInternational$, and $PctIndex$, with relative influences equal to 5.8%, 2.9% and 2.8%, respectively. The BRT model has an in-sample R^2 of 68.7%.

We report the univariate partial dependence plots for the four most important covariates in Subfigures (a) through (d) of Fig. 5. The most relevant covariate, $PctEquityShare$, is nonlinearly related to changes in investors' Sharpe ratios. Investors who have 0% of their invested wealth in equities as self-directed investors experience an improvement in their Sharpe ratio close to 0.3. The relation is very steep until equity levels equal to 10%, and it instead flattens for values between 10% and 40%. Finally, the relation is positive but flat between equity values that range between 40% and 100%. $PctCashShare$ and $PctMutualFunds$ are also nonlinearly related to changes in investors' Sharpe ratios, with effects concentrated on those with 100% of their wealth in cash and 0% of their wealth in mutual funds. Finally, $PctInternational$ is nonlinearly and nonmonotonically related to changes in investors' Sharpe ratios. Those with very little and very large exposures to international markets benefit the most from robo-advice, while those who had between 10% of 40% of their mutual fund wealth in international funds benefit the least. The partial dependence plot for $PctIndex$ (unreported) is negatively sloped, with a convex non-linear shape.

The partial dependence plots highlight that the four most important covariates are non-linearly related to the changes in Sharpe ratios, suggesting that a linear model would be potentially misspecified. To assess this hypothesis, we estimate a kitchen sink linear regression using the same 14 covariates and report which coefficients are statistically significant at the 10% level in the second column of Table 4. The linear regression model explains only 40.1% of the variation in the data, which is 59% of that for BRTs. Furthermore, the kitchen-sink linear regression model estimates as significant, at the 10% level, all regressors except for $Transaction$. Some of these regressors, such as $Male$ and $MgtFees$, are significant because they compensate for the nonlinearities in $PctEquityShare$ and $PctCashShare$. We can see this because they become insignificant when we include a second-order transformation of $PctEquityShare$ and $PctCashShare$ in the third column of Table 4: the number of significant regressors drops from 13 to 11, and the R^2 increases to 55.9% (still very far from the 68.7% obtained by BRTs). If an econometrician were to estimate a kitchen sink regression and did not have access to the results of BRT, however, they would conclude that characteristics such as investors' gender and management fees paid are related to investors' performance improvement when adopting robo-advice, even though the misspecification of the linear model would mostly drive this inference.

In addition, note that the inclusion of the squared term may not properly capture the nonlinear relation between the changes in performance and $PctEquityShare$ that the partial dependence plots uncover. In fact, when we estimate regressions that include higher-order transformations of $PctEquityShare$, we find that these are significant up to the 7th order, suggesting that limiting the analysis to only second-order terms may result in a misspecified model. Once again, an econometrician who wanted to account for the non-linearities across all variables would need to either estimate a kitchen sink regression with all higher-order terms, which would certainly result in an overfit linear model, or perform model selection on the baseline regressors as well as their higher-order terms. In this latter scenario, if they wanted to focus on models that account for up to the 7th order, they would need to estimate an unmanageably large number of models: $2^{14\times7} = 3.169E+29$. Even if they were to obtain a final model that properly captures the nonlinearities, they would likely face great difficulties in interpreting the shape relating each covariate to the dependent variable from the estimated coefficients.

4.3. In- and out-of-sample performance of BRTs

One of the main criticisms against non- and semi-parametric models is that they tend to overfit the training dataset. One could be worried that the non-linearities and non-monotonicities uncovered in Fig. 5 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.

We perform the following Monte Carlo analysis to assess whether BRTs overfit the training dataset. We take the original dataset on which we estimate our BRT results and exclude 1,000 observations. We then estimate the BRT model and test its performance in the holdout sample. 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

[10] 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.

Table 4
Comparing BRTs and linear regression.

Covariate	BRT Relative Influence	Significant in LM-1	Significant in LM-2
PctEquityShare	78.1%	✓	✓
PctCashShare	8.5%	✓	✓
PctMutualFunds	5.8%	✓	✓
PctInternational	2.9%	✓	✓
PctIndex	2.8%	✓	✓
PctStocks	0.6%	✓	✓
Age	0.4%	✓	✓
MgtFees	0.4%	✓	
NumAssets	0.3%	✓	✓
PctETF	0.1%	✓	✓
Tenure	0.1%	✓	✓
Male	0.0%	✓	
Transaction	0.0%		
Volume	0.0%	✓	✓
R-squared	68.7%	40.1%	55.9%

This table reports, in the first column, the relative influence measures for fourteen regressors in explaining the changes in Sharpe ratios pre- and post-advice. Relative influence measures sum to 100% by construction. The second column reports the statistical significance at the 10% level of regression coefficients estimating the same relation (LM-1). The third column reports 10% significance levels in a regression model that includes the squared term for PctEquityShare and PctCashShare in addition to the linear terms (LM-2). The last row of each column reports the R-squared of the associated model. The fourteen covariates included are: Age, the investor's age as of December 2017; $Male$, whether the investor is male; $Tenure$, the tenure of the investor as of December 2017; $NumAssets$, the number of assets held by the investor across accounts; $PctEquityShare$, the percentage of wealth in equities—held directly or through mutual funds; $PctCashShare$, 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; $PctInternational$, the percentage of mutual fund wealth invested in global or international funds—identified using the Lipper mutual fund classification; $MgtFees$ 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 in a month before signing up for advice; $Volume$, the volume (in US dollars) traded by the investors in a month before signing up for advice.

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 (measured as out-of-sample R^2) of the BRT and linear models across all Monte Carlo rounds.

As shown in Section 4.2, the partial dependence plots show that the relation between the regressors and the independent variable is non-linear. As a result, we should expect the linear model to perform relatively poorly compared to BRTs. This is indeed what we find. Subfigure (a) of Fig. 6 shows 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 68%, as shown by the black line. For comparison, note that the linear model has an in-sample R^2 of 40%—green line. The out-of-sample performance of BRTs improves as the number of boosting iterations rises from 100 to approximately 5,000, as shown by the red line. The out-of-sample fit then asymptotes and stabilizes at around 67%, a value greater than the in-sample R^2 of the linear model. The linear model's out-of-sample fit is lower, equaling 39%. As well-known in the literature, adding non-linearities to the linear model by including higher-order transformations of the original regressors improves the in-sample R^2 but deteriorates the out-of-sample R^2 (Friedman et al., 2001).

In Subfigure (b) of Fig. 6, we present the density of the out-of-sample R^2 across simulation rounds. In line with the findings in Subfigure (a), BRTs consistently outperform the linear model out-of-sample.

We summarize this section's findings in *Fact 3*, reported below:

Fact 3*: the investors who benefit the most from robo-advice are the ones who—before adopting advice—did not have high exposure to equities, did not diversify internationally, and did not invest widely in indexed mutual funds.*

5. Beyond performance: Robo-advising and investor attention

So far, we have focused on the investment performance of investors before and after signing up for robo-advice. We now move beyond investment performance and focus on the time spent by investors managing their finances. The optimal inattention models of Abel et al. (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 log into their accounts 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.

We estimate how investors' attention changes after signing up for advice using the 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 $Attention_{i,t}$ is investor's i attention in month t, α_i denote investors' fixed effects, β_t are time-effects, and the dummy variable $ROBO_{i,j,t}$ is equal to zero for every month, except for the jth 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. Finally, the vector $X_{i,t}$ contains standard control variables, such as portfolio return and volatility (Gargano and Rossi, 2018). We report coefficients and 95% confidence intervals based on double-clustered standard errors in Fig. 7.

Fig. 7 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 desktop computers, mobile apps, and mobile browsers, respectively. The number of monthly days with logins increases up to six months after the adoption of advice. Economically, the effect is large for the month when advice is implemented, where individuals, on average, log in five more times. 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. After that, attention becomes more and more negative with the horizon, and the coefficient for the 35th robo-dummy equals -1.5, suggesting

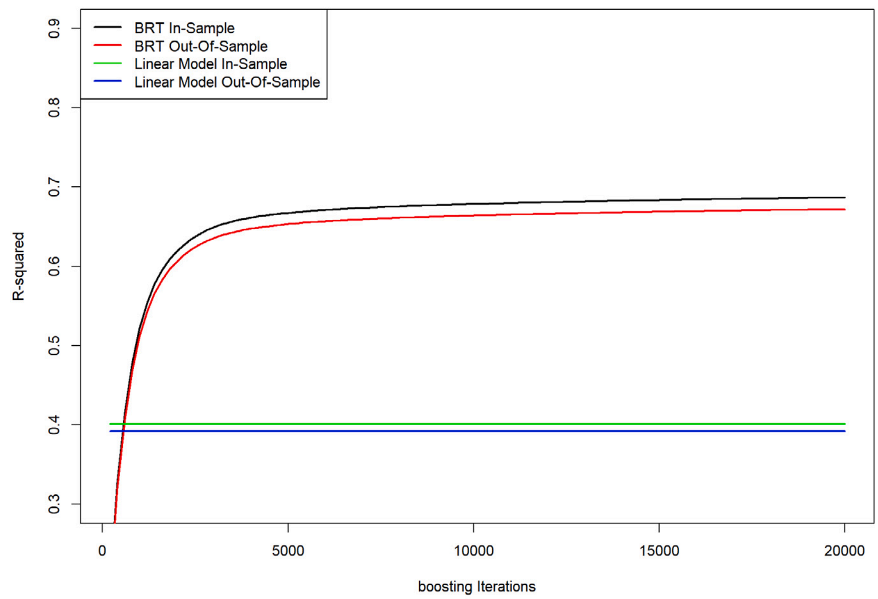

(a) Performance Across Boosting Iterations

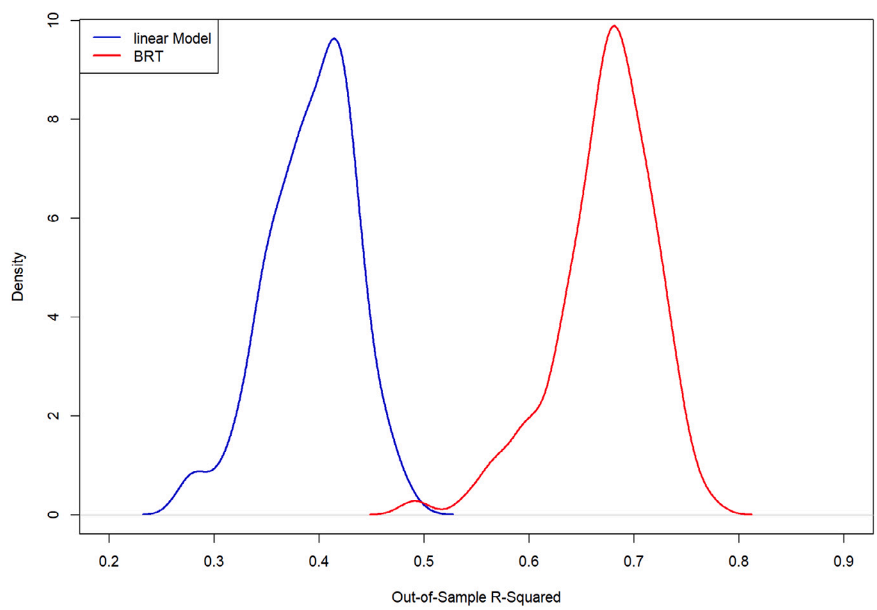

(b) Performance across Monte Carlo Samples

Fig. 6. 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. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

that those who sign up for robo-advice login 1.5 fewer times 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 equals 4.3.

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 where 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 other means decreases and becomes negative over time, suggests two findings. First, after signing up for advice, there is a significant amount of time when individuals log in and actively try to understand what the robo-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 log into their account using

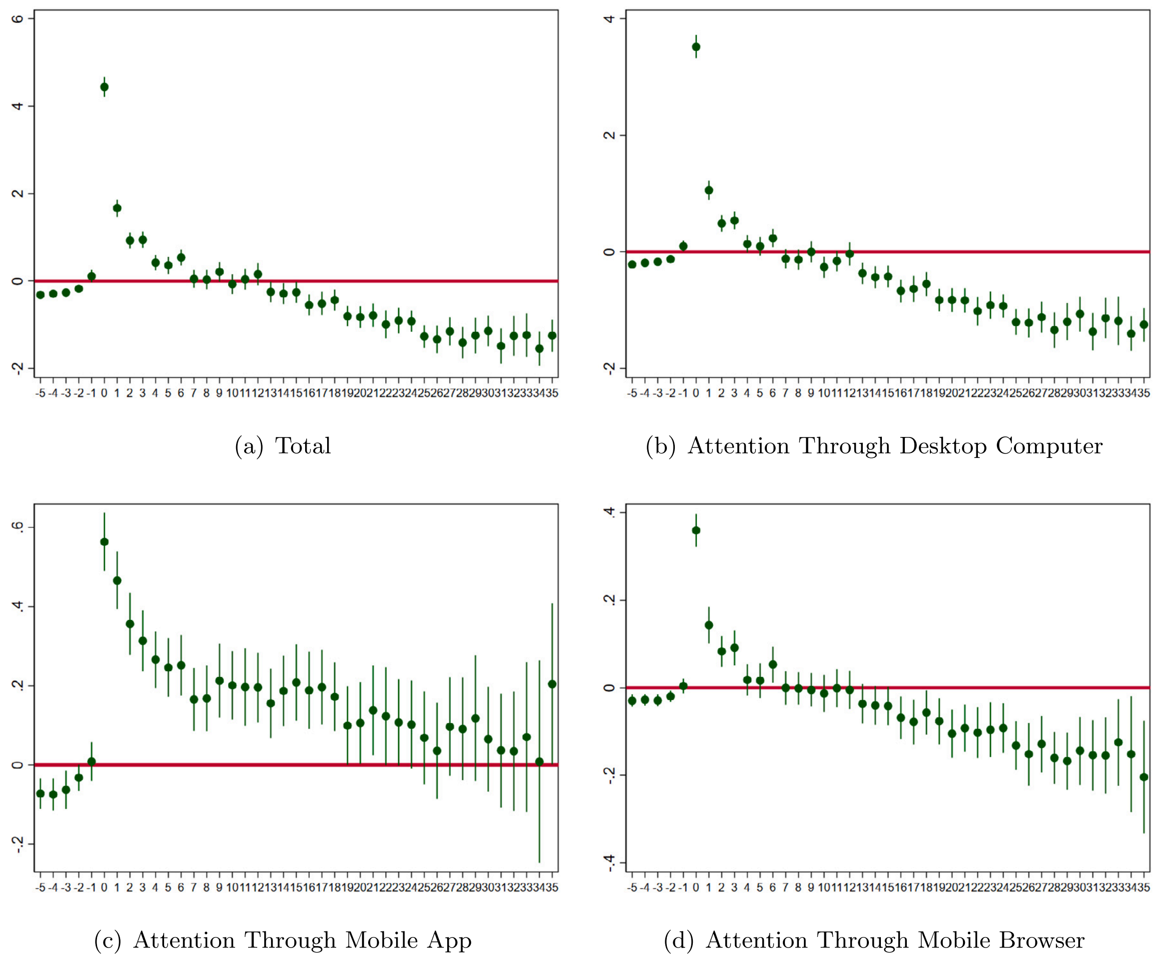

(a) Total

(b) Attention Through Desktop Computer

(c) Attention Through Mobile App

(d) Attention Through Mobile Browser

Fig. 7. Panel regressions relating robo-advising adoption and investors' attention. This figure reports results relating robo-advising adoption and investor attention. 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 jth 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. 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.

the mobile app to gather information regarding the current value of their portfolio.

The results in Fig. 7 use the number of days with logins within a month as a measure of attention. In Figure Online III, 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 login results, with two significant differences. First, as shown in Subfigures (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 min of attention 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 min 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 logins. 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 suggests that, after adopting the service, individuals log in more often but spend less time on the platform each time they log in.

The results reported in the section so far suggest that adopting robo-advice not only improves individuals' investment performance but 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 log in more often to quickly acquire information regarding their portfolio wealth whenever they need to.

We can attach a monetary value to the time saved after adopting robo-advice using investors' wealth and income information. Recall from Table 2 that the median robo-advised investor has an invested wealth of \$494K. From Section 3.2, we also know that the median percentage of wealth invested in our robo-advisor represents 55% of investors' wealth. It follows that the median investor in our sample has a net worth of \$494K/0.55=\$898K. If we assume that our investors are still employed, we can use data from the Survey of Consumer Finances to compute their expected annual income. In particular, know that the ratio between financial assets and income for households in the top decile in terms of net worth—which is where our median investor falls—is \$266.1/\$1.598M=0.166. So we can expect our median investor to have an annual income equal to $0.166 \times \$898K$=\$149.5K and an hourly wage of $\$149.5K/(52 \times 40) = \$72$. From Subfigure (a) of Figure Online III, we know that investors save 30 min per month, or 6 h per year, after adopting robo-advice, so we can conservatively estimate that adopting robo-advice saves investors approximately $\$72 \times 6 = \$431$ per year in terms of their opportunity cost of time.

Finally, combining these results with the diversification effects documented in Facts 1 through 3 suggests that investors not only save time by not trading, but they plausibly avoid investment mistakes that would lead to inefficient investing.

We summarize this section's findings in *Fact 4*, reported below:

Fact 4: *Robo-advised investors spend less time monitoring their portfolios, saving on average six hours of time, valued at roughly \$450 per year.*

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 benefit from both an investment performance perspective and a time-saving perspective. Their benchmark-adjusted performance increases, and the time investors dedicate to their finances decreases. This section analyzes the determinants of adoption and attrition in robo-advice. We aim to understand whether the investors who are likely to benefit more from advice are the ones who (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 who were at—some point—interested in signing up for advice, even though only 31% of them ended up adopting the service. Our data also contain detailed information regarding investors' attrition—that is, which investors decide to stop using the service.

To avoid using machine learning tools in the context of binary outcome variables and survival data and to align our results with other standard results reported in the literature, we proceed with traditional Logit and Cox Proportional Hazard models throughout this section.

6.1. Who signs up for advice?

In the dataset containing all the investors potentially interested in advice, we have 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 enrollment process, the probability rises to 39%, and conditional on enrollment in the robo-advisor, the probability of implementing the robo-advisor's portfolio allocations is 100%.

In what follows, we analyze the determinants of sign-up from the moment individuals start the process of initiating robo-advice adoption. This is when investors are likely to gather precise information about the characteristics of the service. It can, therefore, help us understand which individual characteristics relate to advice adoption. This is crucial because it can help answer whether it is the individuals who need advice the most or those 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 equal 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; $PctETF$ is the client percentage of wealth in ETF; $PctMutualFunds$ is the percentage of wealth in mutual funds; $PctStocks$ is the percentage of wealth in individual stocks; $MgtFees$ are the value-weighted management fees charged by the mutual funds held by the account holders; $PctInternational$ is the percentage of wealth in international stocks or bonds; $Wealth$ is the investor's log invested wealth; $PctEquityShare$ and $PctCashShare$ 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 include appointments categorized as client-facing, individual consults, plan delivery, ad-hoc meetings, check-in, and plan delivery scheduled on the month investors initiated enrollment into robo-advice. All non-binary regressors are standardized so that they have unit variance.

The results are reported in the first column of Table 5. 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 advice process, the more likely they are to sign up.

Concerning the question of whether it is the individuals who benefit the most from advice to be the ones who sign up for it, 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 also likely to sign up for advice if their portfolio is doing poorly, as evidenced by the negative and significant coefficient on $Investor\_Ret$.

Also, individuals holding mutual funds with high management fees, relatively high stocks and cash holdings, and little international diversification are more likely to sign up for robo-advising. Economically, the magnitudes for some of these regressors are rather large. For example, a standard deviation increase in investors' management fees increases the probability of signing up for advice by 3.5 percentage points. Given a constant of 21%, the increase implies approximately a 3.5/21=16.7% increase in the likelihood of signing up for advice.

Other 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 $Wealth$ is particularly large. In the cross-section, a standard deviation increase in wealth lowers the likelihood of signing up for advice by 3.7 percentage points. Once again, economically, the coefficients are very large. Given a constant of 21%, the increase implies approximately a 3.7/21=17.6% reduction in the probability of signing up for advice.

The last variables we consider are clients' interactions with advisors, which relate more to advisors' characteristics than clients'. Investor-scheduled appointments do not seem to impact the probability of signing. What has an impact, on the other hand, are advisor-scheduled appointments. Having an advisor-scheduled appointment on the same month as initiating enrollment into financial advice increases the probability of adopting financial advice by seven percentage points.

In the second column of Table 5, we repeat the analysis using a Logit specification instead of a linear probability model. The results are virtually identical from an economic and a statistical perspective.

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

Table 5
Cross-sectional results on the determinants of robo-advising adoption.

	Linear probability model	Logit model
Market_Ret	−0.002	−0.000
	(−1.25)	(−0.26)
Market_Var	0.022***	0.020***
	(13.98)	(13.54)
Investor_Ret	−0.004**	−0.005***
	(−2.38)	(−3.16)
Realized_Var	0.018***	0.017***
	(11.59)	(11.65)
PctETF	0.024***	0.021***
	(11.90)	(10.81)
PctMutualFunds	0.033***	0.027***
	(6.62)	(5.51)
PctStocks	0.026***	0.023***
	(11.97)	(10.79)
PctInternational	−0.013***	−0.013***
	(−13.01)	(−12.60)
PctEquityShare	−0.032***	−0.031***
	(−19.64)	(−19.32)
PctCashShare	0.020***	0.015***
	(4.60)	(3.48)
MgtFees	0.035***	0.032***
	(34.45)	(32.94)
Wealth	−0.037***	−0.036***
	(−38.68)	(−37.91)
Human_Advisor_Scheduled_Appointments	0.080***	0.075***
	(51.46)	(50.49)
Investor_Scheduled_Appointments	0.004	0.003
	(1.34)	(1.01)
Constant	0.210***	
	(175.73)	
R-squared	0.043	0.037
N	202,571	202,571

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 equal 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; $PctETF$ is the client percentage of wealth in ETF; $PctMutualFunds$ is the percentage of wealth in mutual funds; $PctStocks$ is the percentage of wealth in individual stocks; $MgtFees$ are the value-weighted management fees charged by the mutual funds held by the account holders; $PctInternational$ is the percentage of wealth in international stocks or bonds; $Wealth$ is the investor's log invested wealth; finally, $PctEquityShare$ and $PctCashShare$ 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 $Investor\_Scheduled\_Appointments$ scheduled on the month investors initiated enrollment into robo-advice. These 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. 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.

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 robo-advisory 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 Fig. 8, we present non-parametric estimates of the percentage of investors maintaining their subscription to robo-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 robo-advice. Second, the attrition rate is less than 5% per year. In Subfigure (b), we divide the investors into 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 five percentage points at the four-year mark, suggesting the importance of trust and familiarity with the brand name of the robo-advisor.

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 onboarding 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 who do not. 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 million 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 variation in attrition rates highlights the importance of the human component in robo-advice, because the main

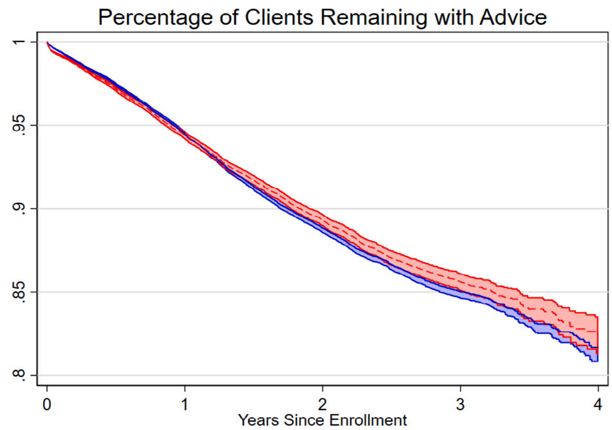

(a) Red: Female Investors; Blue: Male Investors

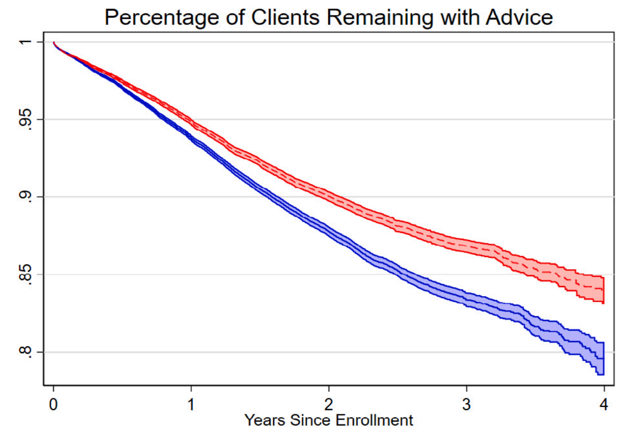

(b) Red: Long-tenure Investors; Blue: Short-tenure Investors

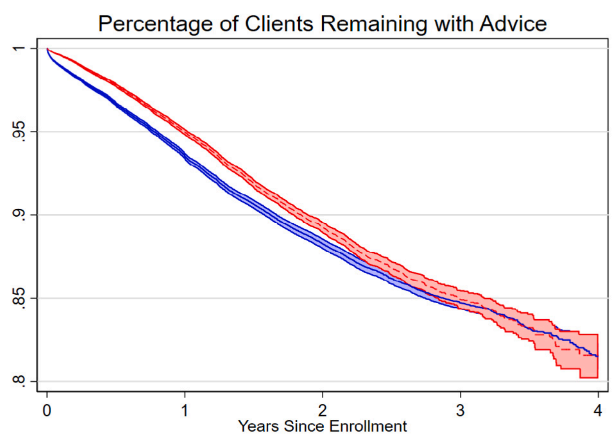

(c) Red: Slow to Enroll; Blue: Quick to Enroll

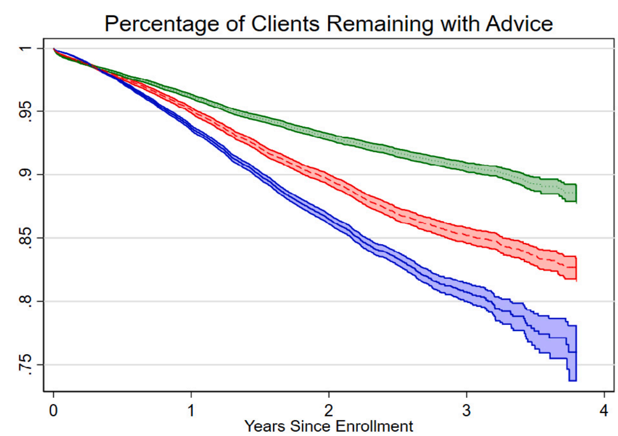

(d) Green: Level 3; Red: Level 2; Blue: Level 1

Fig. 8. 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. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

difference between the three levels of advice stands in the degree of human interactions investors have within the robo-advisor. Level 1 investors do not have a dedicated human advisor, level 2 investors are assigned to a dedicated advisor, and level 3 investors are assigned a dedicated advisor of greater experience and with fewer clients. This different categorization carries significant implications in terms of the interactions investors have with human advisors, as shown in Figure Online IV. The vast majority of level 1 investors interact with a human advisor only on the month they sign up for robo-advice, while level 2 investors have many more interactions, including regular meetings every six months. Level 3 investors have even more interactions with human advisors compared to level 2 investors.

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 6. The first column reports the baseline results, while 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—the hazard ratio equal to 0.912, significant at the 1% level. We find a similar effect for management fees ($MgtFees$). Those paying high management fees before signing up for advice are 5% more reluctant to quit advice after they enroll—hazard ratio equal to 0.96 significant at the 5% 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 13% more likely to quit advice—hazard ratio equal to 1.134, 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, the Level 1 and Level 2 service dummies are highly significant. The base case is Level 3 investors, which are the least

Table 6
Cox-proportional hazard model results on robo-advising attrition.

	Spec 1	Spec 2
Market_Ret	0.991	1.006
	(−0.29)	(0.18)
Market_Var	0.912***	0.923***
	(−3.47)	(−3.00)
Investor_Ret	0.978	0.973
	(−0.75)	(−0.94)
Realized_Var	1.041	1.036
	(1.58)	(1.38)
PctETF	1.021	1.023
	(0.65)	(0.74)
PctMutualFunds	0.862*	0.872*
	(−1.89)	(−1.74)
PctStocks	0.978	0.976
	(−0.65)	(−0.72)
PctInternational	0.997	1.002
	(−0.20)	(0.14)
PctEquityShare	1.134***	1.150***
	(4.00)	(4.39)
PctCashShare	1.033	1.041
	(0.45)	(0.55)
MgtFees	0.956**	0.949***
	(−2.37)	(−2.75)
Wealth	0.991	1.201***
	(−0.52)	(7.68)
Human_Advisor_Scheduled_Appointments	1.052*	1.051*
	(1.82)	(1.77)
Investor_Scheduled_Appointments	0.825***	0.880*
	(−2.66)	(−1.76)
"Level 1" Dummy		2.021***
		(13.73)
"Level 2" Dummy		1.384***
		(6.65)
N	48,993	48,993

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; $PctETF$ is the client percentage of wealth in ETF; $PctMutualFunds$ is the percentage of wealth in mutual funds; $PctStocks$ is the percentage of wealth in individual stocks; $MgtFees$ are the value-weighted management fees charged by the mutual funds held by the account holders; $PctInternational$ is the percentage of wealth in international stocks or bonds; $Wealth$ is the investor's log invested wealth; finally, $PctEquityShare$ and $PctCashShare$ 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 investors have $500K at most in AUM, while Level 2 investors have between $500K and $1M in AUM.

reluctant to quit advice, as shown in Fig. 8. The results in Table 6 show that Level 1 investors are 100% more likely to quit advice than Level 3 investors. Level 2 investors are 40% more likely to quit advice than Level 3 investors.

The results for client wealth highlight the importance of controlling for the type of service investors sign up for. The coefficient on $Wealth$ is insignificantly different from zero in the first column. Once we control for the type of service investors receive, the coefficient on $Wealth$ becomes positive and economically significant. An increase in income increases the probability of quitting advice by 20%—the hazard ratio equals 1.20 and is significant at the 1% level.

Overall, the positive (negative) relation between equity share (management fees) and the probability of quitting advice supports the notion that those investors who benefit the most from advice are the ones who are less likely to leave the service.

We summarize this section's findings in *Fact 5*, reported below:

Fact 5*: The investors who benefit most from the robo-advice service are more likely to adopt it and less likely to quit it.*

7. Conclusions

We study the diversification and welfare effects of a large U.S. hybrid robo-advisor on the portfolios of previously self-directed investors. Across all investors, robo-advising increases indexing and reduces home bias, number of assets held, and fees. These portfolio changes result in superior Sharpe ratios.

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. Our results suggest that the investors who benefit the most from robo-advice are the ones who—before adopting advice—did not have high exposure to equities, did not invest widely in indexed mutual funds, and held a lot of their investable wealth in cash.

We document that the overall time spent making investment decisions decreases after adopting advice. The time savings add up to 6 h per investor, which we value at roughly $450 per year.

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.

Our results suggest that automated portfolio tools can be welfare-enhancing for individual investors who adopt them. An interesting path for future research would be to study the barriers to robo-advising adoption, like algorithmic aversion, i.e., individuals' unwillingness to delegate important decisions to automated algorithms.

CRediT authorship contribution statement

Alberto G. Rossi: Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Stephen Utkus:** Writing – review & editing, Supervision, Project administration, Methodology, Investigation, Conceptualization.

Declaration of competing interest

Alberto Rossi is an unpaid consultant at Vanguard to access the anonymized data. He has no conflicts of interest associated with the results of this work.

Stephen Utkus was employed at Vanguard in a research capacity and retired in August 2020. He has no conflicts of interest associated with the results of this work.

Data availability

Replication Package for "The Diversification and Welfare Effects of Robo-advising" (Reference data) (Mendeley Data).

Appendix A. Supplementary data

Additional tables and figures can be found in the supplementary material online at https://doi.org/10.1016/j.jfineco.2024.103869.

References

Abel, A.B., Eberly, J.C., Panageas, S., 2007. Optimal inattention to the stock market. Amer. Econ. Rev. 92 (2), 244–249.
Abel, A.B., Eberly, J.C., Panageas, S., 2013. Optimal inattention to the stock market with information costs and transactions costs. Econometrica 81 (4), 1455–1481.
Aghion, P., Howitt, P., 1992. A model of growth through creative destruction. Econometrica 60 (2), 323–351.
Badarinza, C., Campbell, J.Y., Ramadorai, T., 2016. International comparative household finance. Annu. Rev. Econ. 8, 111–144.
Barber, B.M., Huang, X., Odean, T., 2016. Which factors matter to investors? Evidence from mutual fund flows. Rev. Financ. Stud. 29 (10), 2600–2642.
Barber, B.M., Odean, T., 2001. Boys will be boys: Gender, overconfidence, and common stock investment. Q. J. Econ. 116 (1), 261–292.
Berk, J.B., Van Binsbergen, J.H., 2016. Assessing asset pricing models using revealed preference. J. Financ. Econ. 119 (1), 1–23.
Bhattacharya, U., Hackethal, A., Kaesler, S., Loos, B., Meyer, S., 2012. Is unbiased financial advice to retail investors sufficient? Answers from a large field study. Rev. Financ. Stud. 25 (4), 975–1032.
Bianchi, M., Briere, M., 2020. Robo-advising for small investors. Available at SSRN 3751620.
Calvet, L.E., Campbell, J.Y., Sodini, P., 2007. Down or out: Assessing the welfare costs of household investment mistakes. J. Polit. Econ. 115 (5), 707–747.
Campbell, J.Y., 2006. Household finance. J. Finance 61 (4), 1553–1604.
Capponi, A., Olafsson, S., Zariphopoulou, T., 2022. Personalized robo-advising: Enhancing investment through client interaction. Manage. Sci. 68 (4), 2485–2512.
Comin, D., Mestieri, M., 2014. Technology diffusion: Measurement, causes, and consequences. In: Aghion, P., Durlauf, S.N. (Eds.), In: Handbook of Economic Growth, vol. 2, Elsevier, pp. 565–622.
D'Acunto, F., 2018a. Identity, beliefs, and choice under risk. Working Paper.
D'Acunto, F., 2018b. Tear down this wall street: Anti-market rhetoric, motivated beliefs, and investment. Working Paper.
D'Acunto, F., Prabhala, N., Rossi, A., 2019a. The promises and pitfalls of robo-advising. Rev. Financ. Stud..
D'Acunto, F., Prokopczuk, M., Weber, M., 2019b. Historical antisemitism, ethnic specialization, and financial development. Rev. Econ. Stud. 86 (3), 1170–1206.
D'Acunto, F., Rossi, A.G., 2021. Robo-advising. In: Palgrave Macmillan Handbook of Technological Finance.
Friedman, J., Hastie, T., Tibshirani, R., 2001. The Elements of Statistical Learning, vol. 1, (no. 10), Springer series in statistics New York, NY, USA.
Gargano, A., Rossi, A.G., 2018. Does it pay to pay attention? Rev. Financ. Stud. 31 (12), 4595–4649.
Gennaioli, N., Shleifer, A., Vishny, R., 2015. Money doctors. J. Finance 70 (1), 91–114.
Gu, S., Kelly, B.T., Xiu, D., 2020. Empirical asset pricing via machine learning. Rev. Financ. Stud. 33 (5), 2223–2273.
Guiso, L., Sapienza, P., Zingales, L., 2008. Trusting the stock market. J. Financ. 63 (6), 2557–2600.
Kim, H.H., Maurer, R., Mitchell, O.S., 2016. Time is money: Rational life cycle inertia and the delegation of investment management. J. Financ. Econ. 121 (2), 427–447.
Linnainmaa, J.T., Melzer, B.T., Previtero, A., 2021. The misguided beliefs of financial advisors. J. Financ. 76 (2), 587–621.
Linnainmaa, J.T., Melzer, B., Previtero, A., Foerster, S., 2018. Financial advisors and risk-taking. Working Paper.
Reher, M., Sokolinski, S., 2023. Robo advisors and access to wealth management. Unpublished Working Paper.
Reher, M., Sun, C., 2019. Automated financial management: Diversification and account size flexibility. J. Invest. Manag. 17 (2), 1–13.
Romer, P.M., 1990. Endogenous technological change. J. Political Econ. 98 (5, Part 2), S71–S102.
Rossi, A.G., Utkus, S., 2019. The needs and wants in financial advice: Human versus robo-advising. Working Paper.