Democratizing Financial Advice*

Michael Reher† and Stanislav Sokolinski‡

August 2019

[Preliminary and Incomplete]

Abstract

We provide evidence that passive asset managers can reduce wealth inequality by lowering a fixed cost of participation in risky asset markets: brokerage account minimums. Using a unique dataset from a large U.S. robo advisor, we study a large and unexpected reduction in the advisor's account minimum. Financially-constrained households respond by significantly increasing the share of their liquid assets invested in the stock market, whereas wealthy households exhibit no response. The shock induces some households to participate in the stock market, thereby increasing the total return on their liquid assets by 2.8 percentage points. These results are consistent with a framework in which households rely on financial advisors to engage in risky asset markets, but brokerage account minimums constrain less-wealthy households' ability to do so. In ongoing work, we use our microeconometric estimates to calibrate a life-cycle model featuring a minimum required investment in risky asset markets, and we will use this model to study how eliminating account minimums affects wealth inequality.

Keywords: FinTech, Stock Market Participation, Life Cycle Models, Inequality

JEL Classification: G11, G24, D14, O3

*We thank John Campbell, Francesco D'Acunto, Ben Friedman, Jakub Jurek, Andrei Shleifer, Celine Sun, Boris Vallee, and seminar participants at Harvard Business School for comments. The views expressed in this paper are those of the author and do not necessarily reflect the position of Wealthfront Inc.

†University of California San Diego, Rady School of Management. Email: mreher@ucsd.edu

‡Rutgers Business School. Email: ssokolinski@business.rutgers.edu

1 Introduction

Recent technological innovation in finance (i.e. "FinTech") bears on a number of long-standing puzzles about households' financial condition (e.g. Philippon 2018). One of the oldest of these puzzles is why households invest so little in the stock market: classic finance theory recommends that all households should at least participate in the stock market − if not take a leveraged position − but fewer than half of U.S. households own stocks at all (e.g. Guiso and Sodini 2013). Over the past decade, automated financial advisors (i.e. "robo advisors") have sought to tap this market of would-be investors, using automation to manage large numbers of portfolios at low per-portfolio marginal cost. Indeed, many leading U.S. robo advisors champion the idea of "democratizing sophisticated financial advice", and their recent growth invites the question of whether automation can increase household participation and investment in the stock market.[1]

We provide evidence that brokerage account minimums constrain household stock market investment, and robo advisors can increase investment by using automation to reduce these minimums. The underlying theory we seek to test begins with the observation that − for a variety of reasons − households rely on financial advisors to invest in the stock market (e.g. Gârleanu and Pedersen 2018; Gennaioli, Shleifer and Vishny 2015; Guiso, Sapienza and Zingales 2008). Traditional advisors incur a per-portfolio cost of management, and so they typically restrict services to portfolios larger than some minimum. Given households' reliance on financial advisors, these account minimums can limit financially-constrained households' allocation to stocks and so constitute a fixed cost of stock market investment (e.g. Vissing-Jørgensen 2002b). Robo advisors, by contrast, can scale portfolio management through automation, enabling them to reduce account minimums and thus increase financially-constrained households' investment in the stock market.

We test this theory in the context of a company-specific experiment in which a leading U.S. robo advisor unexpectedly and significantly lowered its account minimum from $5,000 to $500. This change was motivated by the advisor's philosophy of inclusive investment,

[1]Wealthfront, a leading U.S. robo advisor, describes its founders as seeking to "democratize access to sophisticated financial advice". Similarly, in April 2015 the vice president of operations at Betterment, another leading U.S. robo advisor, wrote that the company's goal was to "democratize sophisticated portfolio management that has traditionally been available only to higher-balance investors".

as well as its hope that less-wealthy households will accumulate enough assets to become highly-profitable customers.[2] We estimate the effects of this shock using a proprietary, household-level dataset with details on a household's liquid assets, investment activity with the robo advisor, demographic information and – for a subset of households – security-level information on their non-robo brokerage accounts. To the best of our knowledge, this is the first paper to use microdata on household investment with a major U.S. robo advisor.

Our main empirical exercise is a difference-in-difference research design in which the "treatment" is the reduction in the robo advisor's account minimum. The "treatment exposure" is the degree to which a household was constrained by the former minimum account size, measured by the size of the reduction relative to the household's liquid assets. We find that the shock disproportionately increased financially-constrained households' allocation to the robo advisor. In particular, households with less than $10,000 in liquid assets – who are overwhelmingly non-participants in the stock market – allocated 59% of their portfolio toward the robo advisor in response to the shock. Consistent with the theory, the magnitude of the reallocation decreases across the wealth distribution, and it equals 11% of the average household's portfolio.

We perform a variety of robustness exercises to evaluate the validity of this research design and the appropriate interpretation of its results. In terms of validity, the results are robust to including multiple fixed effects and heterogeneous time trends by age, income, and geography, which suggests that the estimates are not biased due to spurious correlation with unobserved dynamics that coincide with the shock (e.g. a targeted advertising campaign). In terms of interpretation, the results are similar among a subsample of costly-to-liquidate IRA accounts, suggesting that the baseline results are not driven by substitution across risky positions, but rather by an increase in households' risky share. Neither are the results driven by young households moving forward investments they would have otherwise made later in life (e.g. Fagereng, Gottlieb and Guiso 2017), since we obtain a similar finding among a subsample of relatively-old households.

[2] In the words of the companys then-CEO Adam Nash: "Unlike the many banks and brokerage firms that came before us, [we] refuse to build our business by preying on clients with small accounts. We believe that, given a fair shake, people bold enough to scrape together the savings for their first investment account will build those accounts over time."

To assess the welfare implications of our findings, we follow Calvet, Campbell and Sodini (2007) and calculate the shock's implied effect on a household's total portfolio return, defined as the expected annual return on liquid assets. We find that the shock increases total portfolio return by 2.8 percentage points for households with under $10,000 in liquid assets. Most of these households are non-participants in the stock market, which suggests that removing account minimums can yield welfare gains through stock market participation.

In ongoing work, we use our microeconometric estimates to calibrate a partial equilibrium, life-cycle model featuring a minimum required investment in risky asset markets. We will use this model to study how wealth inequality would change if all financial advisors eliminated account minimums. Given the prevalence of such minimums among financial advisors and our significant microeconometric estimates, it is plausible for such a shock to have meaningful effects through increased stock market participation.

The remainder of the paper proceeds as follows. We conclude this section by situating the paper within the related literature. Section 2 introduces a motivating framework. Section 3 describes the data and natural experiment. Section 4 contains the main analysis, and associated robustness exercises are in Section 5. Section 6 studies welfare implications. Section 7 concludes. All figures and tables are at the end of the main text. The appendix contains additional material.

Related Literature

This paper contributes to three literatures. First, we contribute to a small but growing literature on FinTech. In this vein, the closest papers are D'Acunto, Prabhala and Rossi (2019) and Reher and Sun (2019). These papers study how the introduction of robo advisory products affects diversification and investment mistakes *within* a household's allocation to stocks. By contrast, we focus on how robo advisors can affect the allocation to stocks itself.

Second, we contribute to a literature on financial advice by documenting a previously-unstudied channel through which advisors affect household investment: brokerage account minimums. The prevailing focus in this literature is on how advisors affect households through investment mistakes (e.g. Linnainmaa, Melzer and Previtero 2019; Mullainathan, Noeth and Schoar 2012) and fees (e.g. French 2008; Gil-Bazo and Ruiz-Verdú 2009; Chalmers and Reuter 2018). Our results show how, for a given level of performance and fee structure, financial advisors can affect household portfolios through the account minimums they require.

Third, we contribute to a body of papers focused on limited stock market participation, which spans the asset pricing and household finance literatures. On the asset pricing side, a number of papers have studied how limited participation may contribute to the equity premium puzzle (e.g. Mankiw and Zeldes 1991; Vissing-Jørgensen 2002a; Gomes and Michaelides 2008; Malloy, Moskowitz and Vissing-Jørgensen 2009). A common feature in many of these models is a "fixed cost" of participation in the stock market. We contribute to this literature by identifying a concrete fixed cost that constrains household investment.[3] In so doing, we complement various household finance literatures studying how stock market participation depends on preferences (e.g. Barberis, Huang and Thaler 2006), sophistication (e.g. Grinblatt, Keloharju and Linnainmaa 2011; Christelis, Jappelli and Padula 2010), and education (e.g. Cole, Paulson and Shastry 2014; Van Rooij, Lusardi and Alessie 2011).

2 Motivating Framework

It is theoretically nontrivial that a reduction in account minimums should affect household investment. In a frictionless model, the minimum never binds because households can directly engage the stock market without the assistance of an advisor. Or, even if households require outside assistance, they can borrow to overcome the minimum. To clarify these points, we begin with a simple framework that features two ingredients: (a) inability to invest in risky asset markets without a financial advisor; and (b) limited supply of investible funds. This framework motivates the remaining analysis and will discipline the empirical design in Section 4.

Consider a household solving a static portfolio choice problem. For simplicity, suppose the household has two investment options. First, it can invest in a riskless asset called "cash", which delivers return R^f, where, to minimize notation, $R^f = 0$. Alternatively, it can delegate funds to a financial advisor, who provides a risky, net-of-fee return R.[4] For example, the

[3] Based on a calibration, Haliassos and Bertaut (1995) conclude that account minimums may have a quantitatively small effect on household investment, but they do not actually use data on such minimums.

[4] we take the distribution of returns R as given, so that one can liken this environment to that of a small

household may lack the confidence to invest on its own and thus requires outside assistance to engage with the stock market. Advisors have a cost structure such that they require an account minimum M and charge a management fee of k.[5]

The household has investible assets A and solves the following mean-variance optimization problem,

$$\max_{w} \quad U(w) = \left\{ w\mathbb{E}\left[R\right] - \frac{1}{2}\gamma w^2 \mathrm{Var}\left[R\right] \right\} \tag{1}$$

$$s.t.$$

$$wA \geq M, \tag{2}$$

where w denotes the share of investible assets allocated to the advisor. One can motivate the problem in (1) by supposing that the household has constant absolute risk aversion γ and that R is normally distributed.[6]

Absent the constraint in (2), the household allocates a share

$$\tilde{w} \equiv \frac{1}{\gamma} \frac{\mathbb{E}\left[R\right]}{\mathrm{Var}\left[R\right]}$$

to the financial advisor. However, when the account minimum M is sufficiently large relative to the household's liquid assets, this optimal allocation is no longer feasible. In particular,

open econoour.

[5]A simple way to microfound this outcome is to suppose that advisors are competitive and face the following costs: for portfolios larger than M, the marginal cost per dollar managed is k and there is no fixed cost; for portfolios smaller than M, the marginal cost is still k, but there is a very large fixed cost K, where $K >> M$.

[6]Following Gennaioli, Shleifer and Vishny (2015), the parameter γ may be interpreted as the product of some baseline risk aversion and the household's "anxiety", which depends on who is managing its risky portfolio. We assume the household experiences substantially less anxiety when delegating funds to an advisor relative to investing on its own. Thus, it would never choose to invest in the stock market except through delegation. This assumption is of course a simplification, but it allows me to focus more attention on the role of the account minimum M, which is the focus of this paper.

letting $m \equiv \frac{M}{A}$ denote the relative size of the minimum, the solution to (1) is

$$w^* = \begin{cases} \tilde{w}, & m \leq \tilde{w} \\ m, & m \in (\tilde{w}, 2\tilde{w}) \\ 0, & m \geq 2\tilde{w} \end{cases} . \tag{3}$$

The first condition in equation (3) captures the frictionless case. Here, the constraint (2) does not bind, either because the household is wealthy relative to the account minimum (i.e. m low) or because it already wishes to invest a large share of its wealth with the advisor (i.e. $\tilde{w}$ high). Thus, reducing the minimum has no effect on household investment.

Under the latter two conditions in equation (3), the household is constrained and $w^* \neq \tilde{w}$. These two conditions encode the account minimum's intensive and extensive margin effects, respectively. When the second condition in (3) holds, the household invests the bare minimum, m, so that reducing the minimum actually decreases investment through the intensive margin. Under the third condition, however, reducing the minimum may induce former non-participants to take a risky position, thereby increasing investment through the extensive margin. In Section 4, we find that this third condition is the most empirically-relevant.

3 Data and the Natural Experiment

Our data come from a large U.S. automated financial advisor, Wealthfront, which we will henceforth refer to as the "robo advisor".[7] Wealthfront offers many services including tax loss harvesting, long term financial planning, portfolio lines of credit, and a risk parity fund. Its benchmark product, which is most relevant for this paper, is a portfolio of 10 ETFs across 10 asset classes that is automatically rebalanced.[8] The portfolio weights are determined by a questionnaire which asks the client several questions about her financial

[7]As of March 2018, Wealthfront managed $10 billion and was among the top 5 largest robo advisors in the U.S. market.

[8]Strictly speaking, each asset class has a primary ETF and multiple secondary ETFs. The robo advisor will rebalance toward the secondary ETF if doing so yields a capital loss and thus reduces the client's tax liability.

situation and risk tolerance. The client is assigned to one of 20 possible risk scores, each with its own vector of portfolio weights. Since both taxable and nontaxable accounts (e.g. IRAs) are available, there are a total of 40 possible robo portfolios.

The main dataset contains a weekly time series of client deposits from December 2014 through February 2016. We observe the date and size of the deposit and whether the deposit comes from a new client. We also observe the client's age, annual income, and value of liquid assets, all of which are self-reported via the robo advisor's questionnaire and static. Liquid assets include cash, savings accounts, certificates of deposit, mutual funds, IRAs, 401ks, and public stocks. The upper panel of Table 1 provides summary statistics of the clients in the sample, whom we refer to as "households".

On July 5, 2015, the robo advisor unexpectedly reduced its account minimum from \$5,000 to \$500. This reduction is substantial for many of the households in our sample. For example, a household at the 25^{th} percentile of the wealth distribution (\$45,000) would need to invest only 1% of its liquid assets to set up a robo account, as opposed to 11% under the initial minimum. The reduction did not coincide with a change in the fees charged by the advisor or any other new product launches.[9] Moreover, Section 5 provides evidence that it was not a response to a pre-trend in new deposit formation.

Using the notation from the framework in Section 2, the reduction constitutes a fall in M. In theory, the shock should enable less-wealthy households who were previously constrained by the minimum to invest with the robo advisor. Figure 1 provides stylized evidence consistent with this hypothesis. We plot the empirical wealth distribution for new robo investors before and after the shock date. After the reduction, the wealth distribution shifts left and becomes more concentrated around zero. Building on this intuition, the remainder of this paper estimates the effect of the reduction and assesses its welfare implications.

4 Main Analysis

Our parameter of interest is the average effect of a 1 percentage point reduction in account minimum (relative to a household's investible wealth) on the share of her wealth

[9] The advisor's management fee is 0.25 pps for accounts over \$10,000 and zero for smaller accounts.

invested with the robo advisor. Maintaining the notation from Section 2, this parameter approximately maps to the cross-household average of $\left|\frac{\partial w^*}{\partial m}\right|$, where, again, m is the ratio of account minimum to the household's wealth.

Letting i denote households and t denote weeks, the robo advisor's relative minimum for household i in week t can be written

$$m_{i,t} = \frac{\$5000}{\text{Wealth}_i} - \frac{\$4500}{\text{Wealth}_i} \times \text{Post}_t \equiv \alpha_i - \text{Reduction}_i \times \text{Post}_t, \tag{4}$$

where: Reduction_i is defined as $\frac{\$4500}{\text{Wealth}_i}$; Post_t indicates if t is greater than the week when the reduction occurred (i.e. July 5, 2015); and Wealth_i is the household's investible assets, which theoretically maps to A from the framework in Section 2 and empirically maps to liquid assets as defined in Section 3.

Under an ideal experiment, one would regress the share of the household's total portfolio invested with the robo advisor, henceforth the "robo share", on the regressor in equation (4). In practice, two features of the experiment described in Section 3 prohibit this approach. First, the reduction may have coincided with other dynamics that affect robo investment (e.g. advertising), thus leading to omitted variables bias. It is straightforward to address this bias by including a time fixed effect in the regression or, to account for the possibility of targeted advertising, heterogeneous time trends as we do in Section 5.

Second, we observe a household's liquid assets, which we take as our measure of Wealth_i, but this variable is reported by the household and static. Thus, the naive regression just described would suffer attenuation bias from measurement error. We address this concern by grouping households into 50 quantiles by reported wealth. To gain additional variation, we further partition the data according the household's state of residence. Then, we aggregate households across these two partitions, taking the average across wealth-quantile-by-state bins. We refer to these bins as "cohorts", indexed by j. This procedure is analogous to sorting individual stocks into portfolios when testing asset pricing models.[10] There are

[10] An additional rationale for studying aggregated quantities is that a cohort-level analysis identifies the desired effect using information about cross-sectional units (i.e. cohorts) that both eventually do and do not invest with the robo advisor. By contrast, a household-level analysis would only use information about cross-sectional units (i.e. households) that eventually do invest with the robo advisor, since all households in our data invest at some point.

1,402 such cohorts j, which should be thought of as representative households comprised of the 8,289 individual households i in the microdata.[11] The lower panel of Table 1 provides summary statistics at the cohort level.

Our baseline specification is

$$\Delta\text{Risky Share}_{j,t} = \delta\left(\text{Reduction}_j \times \text{Post}_t\right) + \alpha_j + \alpha_t + u_{j,t}, \tag{5}$$

where, as above: j and t index cohort and week; α_j and α_t are cohort and week fixed effects; and Reduction_j is the reduction in account minimum (i.e. \$4,500) divided by average liquid assets among households in cohort j; and $\Delta\text{Risky Share}_{j,t}$ is the value of deposit flow by cohort j in week t divided by cohort j's liquid assets. We estimate equation (5) over the period from December 7, 2014 through February 28, 2016, so that the event week lies at the midpoint of this window.

The first column of Table 2 contains the results of equation (5). To interpret, a 1 percentage point (pp) reduction in account minimum relative to cohort j's wealth increases the cohort's robo share by 0.02 pps per week. This effect maps to an annualized 59 pp increase in robo share for cohorts in the least-wealthy decile, for which the \$4,500 reduction was equal to 57% of their wealth.[12] By contrast, the corresponding increase in robo share was only 0.3 pp for households in the top decile, for which the reduction was 0.3% of wealth. The table's second column studies the number of deposits, which captures the extensive margin effect characterized in equation (3). The positive point estimate suggests a strong extensive margin response by less-wealthy cohorts. Specifically, a 1 pp reduction in relative account minimum increases the number of deposits by 0.05 log points per week, equal to an annualized 1.6 log points for cohorts in the least-wealthy decile.

Finally, in Figure 2 we study treatment heterogeneity by estimating a non-parametric

[11] Not every wealth quantile is represented in every state. If that were the case, there would instead be $2,500 = 50 \times 50$ cohorts.

[12] Explicitly, $0.59 = 0.02 \times 52 \times 0.57$.

version of equation (5). Specifically, we estimate

$$\Delta\text{Risky Share}_{j,t} = \sum_{d=1}^{9} \delta^d \left(\text{Wealth Decile}_j^d \times \text{Post}_t\right) + \alpha_j + \alpha_t + u_{j,t}, \tag{6}$$

where Wealth Decile$_j^d$ indicates if cohort j belongs to wealth decile d, and the reference category is the 10$^{\text{th}}$ decile. The vertical axis plots the estimated coefficients δ^d. The coefficients may be interpreted as the average increase in robo share for wealth decile d after the reduction, relative to the increase for the 10$^{\text{th}}$ decile. Figure 2 shows how less-wealthy cohorts see a significant increase in robo share, and especially for those with liquid assets below $10,000. As discussed in the next subsection, most of the households in this decile were non-participants in the stock market before the reduction.

4.1 Interpretation

The parameter δ in equation (5) reflects the effect of the reduction on households' robo share. However, to the extent that households finance their robo investment with cash, δ may also be interpreted as the effect on overall risky share. The alternative is that an unobserved risky position was liquidated to finance the robo account. This alternative is unlikely given the costly tax consequences associated with any capital gains from liquidation. Indeed, we obtain similar findings when restricting analysis to costly-to-liquidate retirement portfolios, as discussed in Section 5.3. Thus, the most appropriate interpretation of δ is a mixture of increased risky share and substitution across risky positions.

For a subset of households, the reduction plausibly leads to a switch in stock market participation status. This interpretation is reasonable for households with liquid assets below $10,000. According to the 2013 Survey of Consumer Finances (SCF), only 18% of households with liquid assets below this minimum have either direct or indirect exposure to the stock market.[13] Recall from Figure 2 that such less-wealthy households exhibited the strongest investment response to the reduction. Therefore, roughly 80% of this response comes from

[13]To address concerns that liquid wealth is misreported and thus not appropriate to map to national aggregates, we obtain a similar statistic when considering income, which is arguably less subject to measurement error. Average income for households with less than $10,000 in liquid assets in our data is $57,000, and the stock market participation rate for such households in the 2013 SCF is 26%.

households who were formerly non-participants in the stock market.

5 Robustness and Extensions

5.1 Heterogeneous Time Trends

The week fixed effects α_t in equation (5) account for unobserved dynamics (e.g. advertising) that affect all households equally. However, one might imagine that some households are more exposed to unobserved shocks as, for example, in the case of targeted advertising. To account for this possibility, we reestimate equation (5) after interacting the week fixed effect α_t with the following cohort characteristics: a vector of state indicators; average age; and average income, which also proxies for education. These interaction terms account for the possibility of targeted advertising toward households of a given state, age profile, and income profile, respectively. The resulting estimates for δ in Table 3 are stable across these tests and similar to the baseline estimates in Table 2. It is particularly remarkable that the point estimate is stable when controlling for income-by-week-fixed effects, since income is highly collinear with liquid assets. This finding suggests that the baseline results are not driven by an advertising campaign toward households with low levels of income or education.

5.2 Expected Demand from Less-Wealthy Households

One might also imagine that the robo advisor anticipated strong growth in the number of deposits from less-wealthy households, and it timed the account reduction to amplify this growth. This pre-trend would lead to upwardly biased point estimates. Indeed, recall from Figure 2 that less-wealthy households drive the baseline effect. To check for such a pre-trend, we split the sample of cohorts into a high-wealth and low-wealth group according to median liquid assets across cohorts. Then, we estimate

$$\log\left(\text{New Deposits}_{j,t}\right) = \sum_m \delta^m \left(\text{Below-Median}_j \times \text{Week}_t^m\right) + \alpha_j + \alpha_t + u_{j,t}, \tag{7}$$

where: Week_t^m indicates if week t belongs to month m; and Below-Median_j indicates if cohort j has liquid assets below the median across cohorts.[14] The coefficients δ^m represent the excess new deposit flow by cohorts with below-median wealth in month m, relative to the reference month of June 2015.

Figure 3 plots the results of equation (7). High and low-wealth cohorts display similar behavior in the months leading up to the reduction. However, once the minimum falls, low-wealth households significantly increase their deposit activity. This finding suggests that the baseline estimates in Table 2 are not biased upward due to expectations of short-run demand from less-wealthy investors.

5.3 Irreversible Investment: Retirement Accounts

We now reperform the analysis after filtering out standard taxable accounts, so that the focus is on nontaxable retirement accounts (e.g. IRAs, Roth IRAs). The purpose of this exercise is to assess the extent to which households fund robo investments through cash or liquidation of an unobserved risky position. This distinction matters if the point estimates are to be interpreted as an increase in households' risky share, in contrast to just an increase in robo share.

Suppose households behave according to a mental accounting framework in which they treat their retirement and non-retirement portfolios as segmented. Then households finance their *retirement* robo investment using either cash or liquidated risky positions from a non-robo *retirement* account. Cash is plausibly the dominant source of funds for retirement robo investments, since premature liquidation of a non-robo retirement account would incur a costly penalty. Thus, if the baseline results from Table 2 are borne out among the subsample of retirement accounts, it suggests that the increase in robo share is cash-financed. In this case, the baseline results reflect not only an increase in household robo share, but also an increase in risky share more generally.

The resulting point estimates in Table 4 are larger than the baseline estimates from Table 2. Moreover, they are stable across the heterogeneous time trends introduced in Table

[14] we classify December 2014 and January 2015 as a single month because there was relatively little deposit activity over that period.

3. Given the similarity of the results on this restricted sample, the baseline results likely reflect an increase in risky share, not just a substitution between risky robo and non-robo accounts.

5.4 Testing the Mechanism: Reliance on Financial Advisors

Recall from the framework in Section 2 that account minimums should only affect household investment insofar as households cannot costlessly engage with the stock market on their own. Put differently, the reduction's effect should be stronger where there (a) there are more financial advisors and (b) households trust advisors enough to delegate part of their portfolio to them. We test this hypothesis by reestimating equation (5) after interacting the treatment variable with measures of the market for financial advisors in a given state, all normalized to have zero mean and unit variance.

Table 5 contains the results of this exercise. The interaction in the first column is the number of financial advisors per stock market participant, and it is meant to capture the supply of advisors relative to the base of household investors.[15] The point estimate implies that a 1 standard deviation increase in the supply of advisors increases the effect of the reduction on robo share by 20% (0.04 pps).

The interaction in the second column is the share of financial advisors with a misconduct record from Egan, Matvos and Seru (2019). Higher values of this ratio imply that advisors in a given market are less trustworthy. A 1 standard deviation increase in such untrustworthiness dampens the baseline effect by 20% (0.04 pps). Together with the results from the first column, this result is consistent with the theoretical prediction that account minimums should bind only insofar as households rely on and are willing to trust outside advisors. To be clear, this exercise is not intended to be a rigorous test of a particular model of portfolio delegation. Rather, it lends support to the theory that the effect of account minimums depends on the strength of the relationship between households and advisors.

[15] Our data on the number of financial advisors come from Egan, Matvos and Seru (2019). We measure the number of stock market participants in a state by multiplying the state's population, obtained from the BEA, by the share of tax returns with dividend income, obtained from the IRS. We filter out households with average income less than $50,000 in the IRS data, which corresponds to the 10$^{\text{th}}$ percentile in our data.

6 Welfare Implications

The results to this point do not imply that the reduction in account minimums and subsequent increase in robo share improved household welfare. Indeed, given the well-known underperformance of active managers (e.g. French 2008), it is possible that the reduction actually made households worse off. Mirroring Calvet, Campbell and Sodini (2007), we assess the shock's welfare implications by calculating its effect on a household's total portfolio return, defined as the expected annual return on liquid assets.

Recall from Section 4 that the baseline results imply a 59 pp increase in robo share among households in the least-wealthy decile (i.e. \$10,000), and the effect declines across the wealth distribution to 0.3 pp among the most-wealthy decile (i.e. \$1,300,000). More generally, let Δw_i^* denote the annualized change in robo share for household i due to the reduction in account minimum.[16] As discussed in Section 4.1, the source of this investment can be either cash or a liquidated risky position. In the former case, the increase in the household's total portfolio return is

$$\Delta \bar{R}_{i,\text{Total}} = \Delta w_i^* \times \sigma_{i,\text{Robo}} \times \text{Sharpe}_{i,\text{Robo}} \tag{8}$$

where $\Delta \bar{R}_{i,\text{Total}}$ denotes the change in total portfolio return in excess of the return to cash; $\sigma_{i,\text{Robo}}$ is the total volatility of the household's chosen robo portfolio; and $\text{Sharpe}_{i,\text{Robo}}$ is the robo portfolio's corresponding Sharpe ratio which, as usual, is defined as the ratio of expected excess return to total volatility.

Measuring expected returns from historical data is a well-known challenge, and so we follow Calvet, Campbell and Sodini (2007) and propose an asset pricing model to estimate the expected return for securities in the robo portfolio. Specifically, for each security k, we estimate

$$R_{k,t} = \beta_k^F F_t + \epsilon_{k,t}^F, \tag{9}$$

[16]Explicitly, $\Delta w_i^* = \delta \times 52 \times \frac{4500}{\text{Wealth}_i}$, where δ is the point estimate from column 1 of Table 2 and Wealth_i is the household's liquid assets.

where F_t denotes a column vector of pricing factors in month t; β_k^F denotes the respective row vector of loadings; and $R_{k,t}$ denotes the monthly return on security k in excess of the return to cash and net of expense ratio. The idiosyncratic disturbance $\epsilon_{k,t}^F$ has zero mean and standard deviation σ_k^F.

While imposing a model improves the efficiency of expected return estimates relative to directly measuring them from historical returns, it leads to some bias by imposing an imperfect model of the return structure. Since the choice of model is somewhat arbitrary and the degree of bias will depend on the characteristics of the portfolio in question, we estimate equation (9) separately for five common models indexed by factor vector F. As described in Appendix A.2, these five models are: the standard capital asset pricing model (CAPM), the "market model", the Fama and French three-factor model, a five-factor model augmenting the Fama and French model with global and U.S. bond returns, and a two-factor model based on Vanguard's total equity and bond ETFs. Given the estimated loadings $\hat{\beta}_k^F$ from estimating equation (9) for model F, it is straightforward to compute the expected return on household i's robo portfolio, $\bar{R}_{i,\text{Robo}}^F$. To avoid overweighting any particular model, we consider the average value of $\bar{R}_{i,\text{Robo}}^F$ across models F, denoted $\bar{R}_{i,\text{Robo}}$. This return is net of the robo advisor's 0.25 pps management fee. The robo portfolio's Sharpe ratio is then defined as the ratio of $\bar{R}_{i,\text{Robo}}$ to $\sigma_{i,\text{Robo}}$. The average robo portfolio's Sharpe ratio is 35%, as summarized in Table A1.

Returning to the question of how the robo investment was financed, the alternative to cash financing is that the household liquidates an unobserved risky position. In this case, the increase in the household's total portfolio return becomes

$$\Delta\bar{R}_{i,\text{Total}} = \Delta w_i^* \times \left(\sigma_{i,\text{Robo}} \times \text{Sharpe}_{i,\text{Robo}} - \sigma_{i,\text{Outside}} \times \text{Sharpe}_{i,\text{Outside}}\right), \tag{10}$$

where the notation is the same as in equation (10) with an additional subscript to denote the outside portfolio. Both the outside portfolio's total volatility and its Sharpe ratio are unobserved and must be imputed. For the case of total volatility, the most reasonable approach is to impute a value equal to the total volatility of the chosen robo portfolio.

Consequently, equation (10) simplifies to

$$\Delta \bar{R}_{i,\text{Total}} = \Delta w_i^* \times \sigma_{i,\text{Robo}} \times \left(\text{Sharpe}_{i,\text{Robo}} - \text{Sharpe}_{i,\text{Outside}}\right). \tag{11}$$

We impute the outside portfolio's Sharpe ratio using an auxiliary dataset employed by Reher and Sun (2019). Appendix A.1 describes the dataset in detail. Briefly, it contains a snapshot of households' outside, non-robo portfolios taken in 2016 for a subset of households in our main dataset. The households in the auxiliary sample participated in a program where the robo advisor provided algorithmic financial advice about their outside portfolio. Consequently, for each of the 40 possible robo portfolios, we have security-level information on a set of matched non-robo portfolios. These matched portfolios approximate what an investor interested in robo advising would otherwise be holding in an outside account.

Using the same methodology described above, we estimate the expected return on each non-robo portfolio in the auxiliary dataset, net of management fee. Then, we project the implied Sharpe ratio onto the portfolio holder's observed demographic characteristics: age, log income, and log liquid assets.[17] We use this projection to impute $\text{Sharpe}_{i,\text{Outside}}$ for each household in our main dataset. The average imputed Sharpe ratio is 25%, compared to 35% for the average robo portfolio. Additional summary statistics are in Table A1.

Figure 4 summarizes the results of this exercise. The vertical axis shows the reduction's average effect on total portfolio return for each decile of household liquid assets. The upper bound of each bracket assumes the robo investment was financed by cash as in equation (8). This assumption is more appropriate for less-wealthy households, who, as discussed in Section 4.1, are less likely to have multiple brokerage accounts. In particular, roughly 80% of households in the least-wealthy decile are non-participants in the stock market, per the SCF. For these households, the upper bound provides the most realistic estimate. It suggests that the reduction improves total portfolio return by 2.8 pps. This effect is substantial, and it stems from the fact that the reduction increases these households' robo share by 59 pp on average, much of which reflects a shift from non-participation to participation.

By contrast, the lower bound assumes the investment was financed by a liquidated risky

[17] The coefficients from this regression are -0.05 (0.02), -0.72 (0.17), and 1.10 (0.15) for age, log income, and log liquid assets, respectively. Observations are weighted by portfolio value.

position as in equation (11). The improvement is positive for all households because robo portfolios tend to be relatively well-diversified. However, the effect is negligible for the most-wealthy households. This minor effect does not so much reflect better diversification by the wealthy as it does the fact that they are unconstrained by account minimums, and thus experienced little change in their robo share after the shock. Indeed, if the most-wealthy decile invested with the same efficiency as the robo advisor, the return on its stock portfolio would improve by an imputed 0.96 pps due to better diversification.[18]

7 Conclusion

We found that brokerage account minimums are a constraint on household investment in the stock market, and automation can ease this constraint by enabling financial advisors to manage large numbers of arbitrarily small portfolios. We arrived at this conclusion by studying a natural experiment where a large U.S. robo advisor suddenly reduced its account minimum by a factor of 10. The shock disproportionately increased investment by less-wealthy households, who were constrained by the initial minimum. In particular, the results appear to be driven by a subset of households who became stock market participants as a consequence of the reduction. For such households, access to financial advice increased their expected return on liquid assets by 2.8 pps.

These findings exemplify how advancements in "FinTech" can enable more households to reap the benefits of access to financial markets. In ongoing work, we calibrate a partial equilibrium, life-cycle model featuring a minimum required investment in risky asset markets. We will use this model to study how stock market participation and wealth inequality would change if all financial advisors eliminated account minimums. Our current, microeconometric estimates suggest that such a shock could plausibly have meaningful effects.

[18] Explicitly, the average imputed value of $\sigma_{i,\text{Robo}} \times \left(\text{Sharpe}_{i,\text{Robo}} - \text{Sharpe}_{i,\text{Outside}}\right)$ is 0.71 for the most-wealthy decile, to which we add back the robo advisor's 0.25 pp management fee.

References

Barberis, N., Huang, M. and Thaler, R. H.: 2006, Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing, *American Economic Review* .

Calvet, L. E., Campbell, J. and Sodini, P.: 2007, Down or Out: Assessing the Welfare Costs of Household Investment Mistakes., *Journal of Political Economy* .

Chalmers, J. and Reuter, J.: 2018, Is Conflicted Investment Advice Better than No Advice?, *NBER Working Paper* .

Christelis, D., Jappelli, T. and Padula, M.: 2010, Cognitive Abilities and Portfolio Choice, *European Economic Review* .

Cole, S., Paulson, A. and Shastry, G. K.: 2014, Smart Money? The Effect of Education on Financial Outcomes, *Review of Financial Studies* .

D'Acunto, F., Prabhala, N. and Rossi, A. G.: 2019, The Promises and Pitfalls of Robo-advising, *Review of Financial Studies (forthcoming)* .

Egan, M., Matvos, G. and Seru, A.: 2019, The Market for Financial Adviser Misconduct, *Journal of Political Economy (forthcoming)* .

Fagereng, A., Gottlieb, C. and Guiso, L.: 2017, Asset Market Participation and Portfolio Choice over the Life Cycle, *The Journal of Finance* .

French, K. R.: 2008, Presidential Address: The Cost of Active Investing, *The Journal of Finance* **63**.

Gârleanu, N. and Pedersen, L. H.: 2018, Efficiently Inefficient Markets for Assets and Asset Management, *The Journal of Finance* .

Gennaioli, N., Shleifer, A. and Vishny, R.: 2015, Money Doctors, *Journal of Finance* **70**(1).

Gil-Bazo, J. and Ruiz-Verdú, P.: 2009, The Relation between Price and Performance in the Mutual Fund Industry, *The Journal of Finance* .

Gomes, F. and Michaelides, A.: 2008, Asset Pricing with Limited Risk Sharing and Heterogeneous Agents, *Review of Financial Studies* .

Grinblatt, M., Keloharju, M. and Linnainmaa, J.: 2011, IQ and Stock Market Participation, *The Journal of Finance* .

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

Guiso, L. and Sodini, P.: 2013, Household Finance: An Emerging Field, *Handbook of the Economics of Finance*, Vol. 2, Elsevier.

Haliassos, M. and Bertaut, C. C.: 1995, Why do so Few Hold Stocks?, *The Economic Journal* .

Linnainmaa, J. T., Melzer, B. T. and Previtero, A.: 2019, The Misguided Beliefs of Financial Advisors, *The Journal of Finance (forthcoming)* .

Malloy, C. J., Moskowitz, T. J. and Vissing-Jørgensen, A.: 2009, Long-Run Stockholder Consumption Risk and Asset Returns, *The Journal of Finance* .

Mankiw, G. N. and Zeldes, S. P.: 1991, The Consumption of Stockholders and Nonstockholders, *Journal of Financial Economics* .

Mullainathan, S., Noeth, M. and Schoar, A.: 2012, The Market for Financial Advice: An Audit Study, *NBER Working Paper* .

Philippon, T.: 2018, The FinTech Opportunity, *NBER Working Paper* .

Reher, M. and Sun, C.: 2019, Automated Financial Management: Diversification and Clientele Effects, *Journal of Investment Management* .

Van Rooij, M., Lusardi, A. and Alessie, R.: 2011, Financial Literacy and Stock Market Participation, *Journal of Financial Economics* .

Vissing-Jørgensen, A.: 2002a, Limited Asset Market Participation and the Elasticity of Intertemporal Substitution, *Journal of Political Economy* .

Vissing-Jørgensen, A.: 2002b, Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structure, *NBER Working Paper* .

Figures

Figure 1: Wealth Distribution Around Reduction in Account Minimum

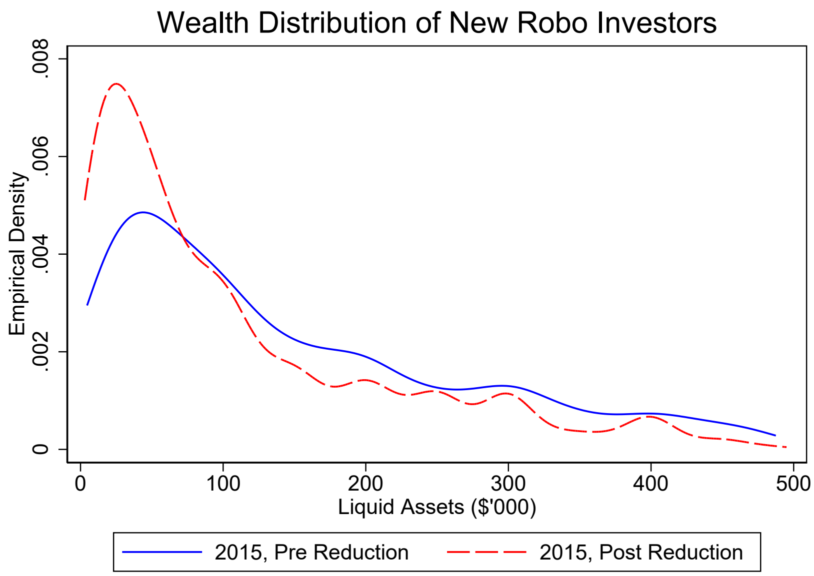

Note: This figure plots the empirical density of household liquid assets for new investors with the robo advisor. The blue solid curve corresponds to the period from January 1, 2015 through July 5, 2015, and the red dashed curve corresponds to the remaining period in 2015. The plot excludes households with liquid assets above $500,000. The density is based on a Gaussian kernel.

Figure 2: Effects of Reduction in Account Minimum by Wealth Decile

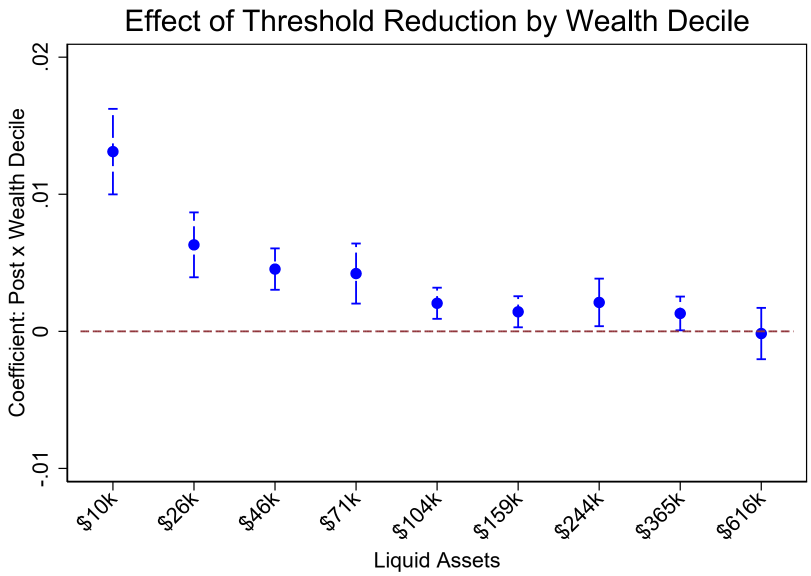

Note: This figure plots results from regressions of the form $\Delta\text{Risky Share}_{j,t} = \sum_{d=1}^{9} \delta^d \left(\text{Wealth Decile}_j^d \times \text{Post}_t\right) + \alpha_j + \alpha_t + u_{j,t}$, where subscripts j and t denote cohort and week; Wealth Decile$_j^d$ indicates if cohort j belongs to wealth decile d; α_j and α_t are cohort and week fixed effects; and the remaining notation and sample period are the same as in Table 2. The vertical axis plots the estimate of δ^d for the first 9 deciles, since the 10^{th} decile is the reference category. The horizontal axis shows the average liquid assets for cohorts in each decile. The brackets are a 95% confidence interval based on standard errors clustered by cohort.

Figure 3: Deposit Flow by Above and Below-Median Wealth Households

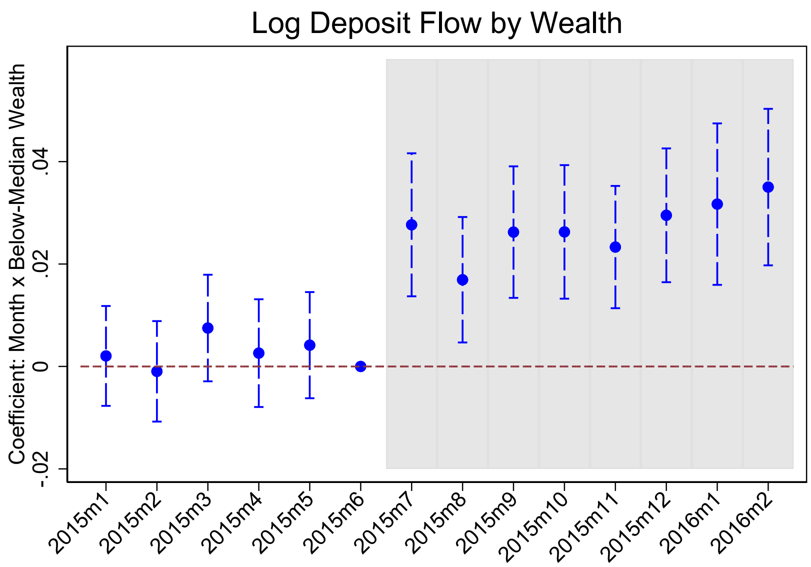

Note: This figure plots results from regressions of the form $\log\left(\text{New Deposits}_{j,t}\right) = \sum_m \delta^m \left(\text{Below-Median}_j \times \text{Week}_t^m\right) + \alpha_j + \alpha_t + u_{j,t}$, where subscripts j and t denote cohort and week; Week_t^m indicates if week t belongs to month m; Below-Median_j indicates if cohort j has liquid assets below the median across cohorts; α_j and α_t are cohort and week fixed effects; and the remaining notation and sample period are the same as in Table 2. The vertical axis plots the estimate of δ^m, where the reference month is June 2015. The shaded region corresponds to the period after which the reduction in account minimum occurred. The brackets are a 95% confidence interval based on standard errors clustered by cohort.

Figure 4: Effect on Household Total Portfolio Return by Wealth Decile

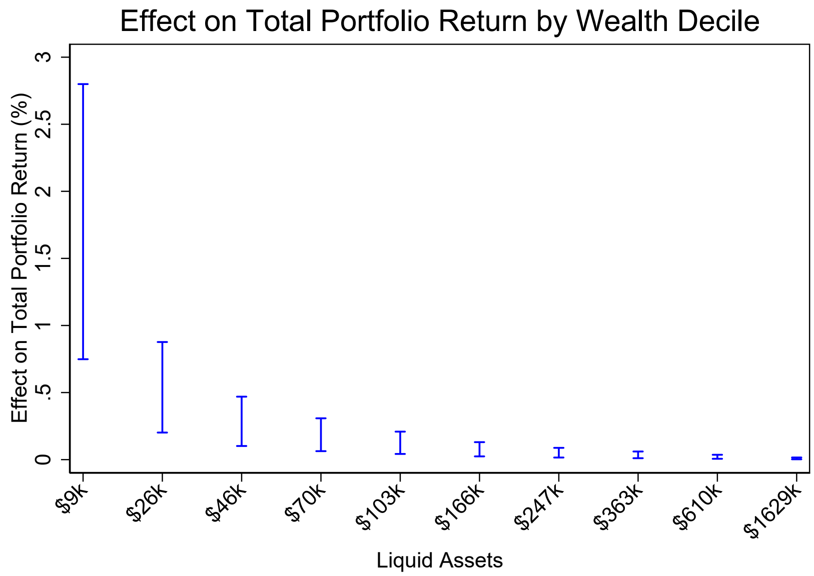

Note: This figure plots the estimated effect of the reduction in account minimum on household total portfolio return by wealth decile. Total portfolio return is defined as the return on liquid assets. The upper end of each bracket assumes the robo investment was financed by cash, and the bottom end assumes it was financed by a liquidated risky position. In the latter case, the outside portfolio's total volatility is imputed as the volatility of the chosen robo portfolio, and its Sharpe ratio is imputed based on the household's age, log income, and log liquid assets and auxiliary information about household investment behavior from Reher and Sun (2019). See Section 6 for additional details.

Tables

Table 1: Summary Statistics

	Mean	Standard Deviation	Percentiles 25^{th}	50^{th}	75^{th}
Household Level:					
Age_i	35.25	9.31	29	33	40
Income_i (\$1,000)	130.33	100.86	69	100	160
Liquid Assets_i (\$1,000)	324.07	546.97	45	123	350
Number of Households: 8,289					
Number of Household-Weeks: 505,629					
Cohort Level:					
Age_j	36.54	8.48	30.41	34.82	40.71
Income_j (\$1,000)	124.55	83.77	70	102.58	151.01
Liquid Assets_j (\$1,000)	342.39	573.04	40.74	120	396.2
Reduction_j (%)	11.28	20.12	1.14	3.75	11.05
ΔRisky $\text{Share}_{j,t}$ (%)	0.52	4.58	0.01	0.01	0.05
Advisors-per-Investor_j (%)	0.91	0.45	0.58	0.8	1.2
Advisors with Misconduct_j (%)	7.87	1.67	6.53	7.52	9.06
Number of Cohorts: 1,402					
Number of Cohort-Weeks: 85,522					

Note: This table presents summary statistics of the main dataset. Subscripts i, j, and t denote household, cohort, and year. The upper panel summarizes household-level variables. Liquid Assets include cash, savings accounts, certificates of deposit, mutual funds, IRAs, 401k plans, and public stocks. The lower panel summarizes cohort-level variables. Cohorts are aggregates of households in wealth-by-state bins, as described in Section 4. Reduction_j is the value of the reduction in account minimum (i.e. \$4,500) divided by the average liquid assets for households in cohort j. ΔRisky $\text{Share}_{j,t}$ is the average robo investment by cohort j in week t divided by cohort j's average liquid assets. Advisors-per-Investor_j is the number of financial advisors divided by the number of stock market participants in 2015, measured using the share of households with dividend income based on IRS data, in j's state of residence. Advisors with Misconduct_j is the share of financial advisors with misconduct records from Egan, Matvos and Seru (2019) in j's state of residence. Observations in the upper panel are household-weeks. Observations in the lower panel are cohort-weeks. The sample period is December 7, 2014 through February 28, 2016, and there are 61 weeks in the sample.

Table 2: Effect of Account Minimums on Household Investment

Outcome	ΔRisky Share$_{j,t}$	$\log\left(\text{New Deposits}_{j,t}\right)$
Reduction$_j$ $\times$ Post$_t$	0.020^{**}	0.054^{**}
	(0.003)	(0.010)
Week FE	Yes	Yes
Cohort FE	Yes	Yes
R-squared	0.030	0.216
Number of Observations	85522	85522

Note: Subscripts j and t denote cohort and week. This table estimates equation (5). Reduction$_j$ is the value of the reduction in account minimum (i.e. $4,500) divided by the average liquid assets for households in cohort j. Post$_t$ indicates if t is greater than the week of July 5, 2015. ΔRisky Share$_{j,t}$ is the average robo investment by cohort j in week t divided by cohort j's average liquid assets. New Deposits$_{j,t}$ is the number of deposits by cohort j in week t. Observations are cohort-weeks. Cohorts are aggregates of households in wealth-by-state bins, as described in Section 4. The sample period is December 7, 2014 through February 28, 2016. Standard errors clustered by cohort are in parentheses.

Table 3: Robustness to Non-Parametric Heterogeneous Time Trends

Outcome	ΔRisky $\text{Share}_{j,t}$		
$\text{Reduction}_j \times \text{Post}_t$	0.020**	0.020**	0.020**
	(0.003)	(0.003)	(0.003)
Week FE	No	Yes	Yes
State-Week FE	Yes	No	No
Age-Week FE	No	Yes	No
Income-Week FE	No	No	Yes
Cohort FE	Yes	Yes	Yes
R-squared	0.067	0.031	0.031
Number of Observations	85400	85522	85522

Note: Subscripts j and t denote cohort and week. This table estimates equation (5) with the inclusion of additional time-varying fixed effects. Column 1 includes state-by-week fixed effects. Columns 2 and 3 interact week fixed effects with the average age and log income for households in cohort j, respectively. The remaining notation is the same as in Table 2. Observations are cohort-weeks. Cohorts are aggregates of households in wealth-by-state bins, as described in Section 4. The sample period is December 7, 2014 through February 28, 2016. Standard errors clustered by cohort are in parentheses.

Table 4: Robustness to Subsample of Nontaxable Accounts

Outcome	ΔRisky Share$_{j,t}$			
Reduction$_j$ $\times$ Post$_t$	0.066**	0.066**	0.069**	0.063**
	(0.028)	(0.028)	(0.029)	(0.029)
Week FE	Yes	No	Yes	Yes
State-Week FE	No	Yes	No	No
Age-Week FE	No	No	Yes	No
Income-Week FE	No	No	No	Yes
Cohort FE	Yes	Yes	Yes	Yes
R-squared	0.027	0.082	0.029	0.030
Number of Observations	40138	39894	40138	40138

Note: Subscripts j and t denote cohort and week. This table estimates equation (5) on the subsample of nontaxable accounts. Columns 2-4 include the additional time varying fixed effects from Table 3. The remaining notation is the same as in Table 2. Observations are cohort-weeks. Cohorts are aggregates of households in wealth-by-state bins, as described in Section 4. The sample period is December 7, 2014 through February 28, 2016. Standard errors clustered by cohort are in parentheses.

Table 5: Heterogeneity by Market for Financial Advisors

Outcome	ΔRisky Share$_{j,t}$	
Reduction$_j$ $\times$ Post$_t$	0.020**	0.021**
	(0.003)	(0.003)
Reduction$_j$ $\times$ Post$_t$ $\times$ Interaction$_{s(j)}$	0.004*	-0.004*
	(0.002)	(0.002)
Interaction	Advisors per Investor	Advisors with Misconduct
Interaction-Week FE	Yes	Yes
Week FE	Yes	Yes
Cohort FE	Yes	Yes
R-squared	0.030	0.031
Number of Observations	80764	82594

Note: Subscripts j and t denote cohort and week. This table estimates equation (5) with the inclusion of interaction terms that describe the market for financial advisors in the state $s(j)$ where cohort j resides. The interaction in column 1 is the number of financial advisors divided by the number of stock market participants in 2015, measured using the share of households with dividend income based on IRS data. The interaction in column 2 is the share of financial advisors with misconduct records from Egan, Matvos and Seru (2019). All interactions are normalized to have zero mean and unit variance. Interaction-Week FE denotes the product of the interaction variable with a vector of week fixed effects. The remaining notation is the same as in Table 2. Observations are cohort-weeks. Cohorts are aggregates of households in wealth-by-state bins, as described in Section 4. The sample period is December 7, 2014 through February 28, 2016. Standard errors clustered by cohort are in parentheses.

A Appendix

A.1 Additional Data Description

This section describes the auxiliary dataset mentioned in Section 6, which comes from Reher and Sun (2019). It contains snapshots of households' portfolio holdings in an outside, non-robo brokerage account. The dataset was generated by a free online tool through which the data provider gave financial advice to clients about their outside portfolio holdings. Specifically, clients would provide their log-in credentials for their outside account. Then, the robo advisor would take a snapshot of the account holdings and run an advice-generating algorithm on it. This produces a set of snapshots of households' non-robo accounts. While the advice algorithm ran, the robo advisor would ask the household to answer its standard questionnaire meant to gauge risk preferences, which is the source of the demographic variables used in this paper. Finally, at the conclusion of the report, the advisor would provide an unbiased analysis of the household's non-robo and robo portfolios. Thus, we observe a matched non-robo portfolio for each household in the sample. The tool was launched in January 2016, and our sample contains 1,180 household-level snapshots taken between January 2016 and November 2016. We merge this dataset with security level information from CRSP to produce a cross-sectional dataset of households' robo and non-robo portfolios.

A.2 Estimating Portfolio Returns

Given the estimated loadings $\hat{\beta}_k^F$ and idiosyncratic volatilities $\hat{\sigma}_k^F$ from estimating equation (9) for model F, we compute the estimated mean μ_p^F and variance $\left(\sigma_p^F\right)^2$ of excess returns on portfolio p as

$$\mu_p^F = \left(\sum_k w_{k,p} \hat{\beta}_k^F\right) \mu^F \tag{A1}$$

$$\left(\sigma_p^F\right)^2 = \sum_k \left(w_{k,p} \hat{\sigma}_k^F\right)^2 + \left(\sum_k w_{k,p} \hat{\beta}_k^F\right) \Sigma^F \left(\sum_k w_{k,p} \hat{\beta}_k^F\right)', \tag{A2}$$

where μ^F is the expected value of F_t, Σ^F is the covariance matrix of F_t, and $w_{k,p}$ is the weight of security k in portfolio p. The weights $w_{k,p}$ are based on the subportfolio consisting of stocks, mutual funds, and exchange traded funds (ETFs). This is because bonds and options are held by few portfolios in the sample, and pricing them is less straightforward (Calvet, Campbell and Sodini 2007).

we estimate equation (9) for the following five asset pricing models,

$$F_t^{CAPM} = [R_t^m]', \tag{A3}$$

$$F_t^{MKT} = [R_t^m, \quad 1]', \tag{A4}$$

$$F_t^{FF} = \left[R_t^m, \quad R_t^{HML}, \quad R_t^{SMB}\right]', \tag{A5}$$

$$F_t^{FF+} = \left[R_t^m, \quad R_t^{HML}, \quad R_t^{SMB}, \quad R_t^{USB}, \quad R_t^{GLB} \quad \right]', \tag{A6}$$

$$F_t^{VAN} = \left[R_t^{VT}, \quad R_t^{BND}\right]', \tag{A7}$$

$$\tag{A8}$$

where R_t^m is the monthly market return based on the global Morgan Stanley Capital International Index (MSCII); R_t^{HML} is the monthly return between high book-to-market stocks and low book-to-market stocks; R_t^{SMB} denotes the spread in monthly return between stocks with a small market capitalization and a big market capitalization; R_t^{USB} is the monthly return on the Barclays Aggregate U.S. Bond Index Unhedged; R_t^{GLB} is the monthly return on the Barclays Global Aggregate Bond Index Unhedged; R_t^{VT} is the monthly return on Vanguard's total stock market ETF (VT); and R_t^{BND} denotes the return on Vanguard's total bond market ETF (BND). The returns $R_t^m, R^{VT}, R_t^{BND} R_t^{USB}$, and R_t^{GLB} are in excess of the return to cash, which we measure as the one month Treasury yield.

In words, equations (A3)-(A7) are: the standard capital asset pricing model (CAPM), the "market model", the Fama and French three-factor model, a five-factor model augmenting the Fama and French model with global and U.S. bond returns, and a two-factor model based on Vanguard's total equity and bond ETFs. Our data on monthly returns come from the Center for Research in Security Prices (CRSP) and Kenneth French's website. We use the sample mean and covariance matrices to calibrate the moments of the factors. For the CAPM, these are $\mu^{CAPM} = 0.068$ and $\Sigma^{CAPM} = 0.170$. To obtain annualized estimates, we multiply the estimated mean and variance from equations (A1)-(A2) by 12.

In terms of data cleaning, we winsorize the sample according to the estimated moments from equations (A1)-(A2) by 2.5% on both sides. We also drop brokerage portfolios under $100 in value. We use the longest available time series of monthly returns for each security $R_{k,t}$ and factor F_t dating back to January 1975.

Appendix Tables and Figures

Table A1: Characteristics of Robo and Matched Non-Robo Portfolios

	Mean	Standard Deviation	Percentiles 25th	50th	75th
Sharpe Ratio (%):					
Robo	34.98	5.51	29.95	37.12	39.28
Non-Robo	24.93	8.31	18.08	26.2	31.55
Expected Return (%):					
Robo	5.11	.61	4.61	4.78	5.7
Non-Robo	4.76	1.17	4.1	4.73	5.39
Idiosyncratic Variance: (% Total Variance):					
Robo	15.58	2.5	13.66	16.16	17.07
Non-Robo	46.6	25.92	22.29	43.42	72.16
Number of Households: 1,180					

Note: This table presents summary statistics of the auxiliary dataset described in Appendix A.1. Observations are households who participated in a program where the robo advisor provided algorithmic financial advice about their non-robo portfolio. The portfolio characteristics indicated in the leftmost column are based on a factor model and estimated using the methodology described in Section 6. For each characteristic, we calculate summary statistics for the household's robo portfolio and its matched non-robo portfolio.