Robo Advisors and Access to Wealth Management*

Michael Reher† and Stanislav Sokolinski‡

October 2022

Abstract

We examine how access to automated wealth managers affects household investment and welfare across the wealth distribution. Our setting features novel microdata from a major U.S. robo advisor and a quasi-experiment in which the advisor reduces its account minimum by 90%. Based on a difference-in-difference estimator, the reduction significantly increases middle-class households' participation but does not affect wealthier or poorer households. We rationalize this behavior with a quantitative model calibrated using portfolio-level data. The reduction raises welfare through diversification, priced risk exposure, and personalization. The overall welfare gains are moderate but heterogeneous. Older households with weak earnings growth gain most.

Keywords: FinTech, Financial Advice, Portfolio Delegation, Inequality

JEL Classification: G11, G24, D3, O3

*We thank Christoph Aymanns, John Campbell, Francesco D'Acunto, Ben Friedman, Pedro Gete, Steffen Hitzemann, Jakub Jurek, Iva Kalcheva, Yaron Levi, Jun Liu, Matteo Maggiori, Apurva Mehta, Will Mullins, Alessandro Previtero, Franklin Qian, Enrichetta Ravina, Jonathan Reuter, Alberto Rossi, Andrei Shleifer, Celine Sun, Ross Valkanov, Boris Vallee, and seminar participants at Northeastern University Finance Conference, University of Mannheim, USC, University of Wisconsin, Jackson Hole Winter Finance Conference, FINRA NORC Diversity Conference, FMA Annual Conference, Paris FinTech Conference, Toronto FinTech Conference, the New York Fed's FinTech Conference, Georgetown's Finance and Policy Center, UC San Diego, Rutgers Business School, California Corporate Finance Conference, CAFR FinTech Workshop, and Harvard's finance PhD lunch for comments. Ruikai Ji provided excellent research assistance. The views expressed in this paper are those of the authors 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, Newark and New Brunswick. Email: ssokolinski@business.rutgers.edu

1 Introduction

Recent technological innovation has scaled the provision of traditional financial services (e.g., Philippon (2016)). The wealth management market, for example, has seen rapid growth in the market share of "robo advisors": intermediaries that use automation to provide services similar to a traditional wealth manager (D'Acunto and Rossi (2020)). Robo advisors have grown exponentially over the past decade, and all of the Big-4 U.S. banks now offer a robo advising service.[1]

Automation may improve the provision of wealth management in two ways. First, it may help households who already have access to professional management to invest more efficiently. For example, automation may limit the scope for impulsive recommendations, conflicts of interest, and various other inefficiencies in the market for financial advice.[2] Second, automation may enable robo advisors to extend these efficiency gains to households who cannot access professional management by lowering costs of having a wealth manager. This conjecture has an intuitive appeal and generated substantial media attention.[3] However, testing it empirically has proven elusive.

We examine how robo advisors affect less-wealthy households. We do so through a quasi-experiment that expands access to automated wealth management, and, subsequently, through a quantitative model that helps understand why the less-wealthy seek wealth management in the first place. We have three main findings. First, the expansion sharply and significantly raises participation in automated wealth management by the middle class. This finding offers a rare example of non-affluent households participating in wealth management. It challenges both the canonical model of a self-sufficient investor as well as models in which households do not seek outside management because of, say, mistrust (e.g., Gennaioli, Shleifer and Vishny (2015))) or financial illiteracy (e.g., Van Rooij, Lusardi and Alessie (2011)). Second, robo-managed portfolios feature greater diversification and exposure to priced risk factors than portfolios managed by households themselves. Embedding these portfolio differences into a canonical model of portfolio choice can rationalize the empirical effect. Third, access to automated wealth management yields moderate

[1] The top five robo advisors managed $283 billion in 2020 versus $30.4 billion in 2015 (Appendix Table A1). The Big-4 banks are Bank of America, Citigroup, JP Morgan, and Wells Fargo.

[2] These inefficiencies include impulsive recommendations (e.g., Linnainmaa, Melzer and Previtero (2021)), pandering to client biases (e.g., Mullainathan, Noeth and Schoar (2012)), overly generic recommendations (e.g., Foerster et al. (2017)), and recommendations of inferior products due to commissions (e.g., Egan (2019); Chalmers and Reuter (2020)). D'Acunto and Rossi (2020) discuss how robo advisors can correct some of these inefficiencies.

[3] "The wealth-management industry stratifies customers in a manner rather similar to airlines. High-net-worth clients fly business class, picking stocks and chatting in person with named advisors. Cattle class gets no service at all. Technology is conspiring to change that" (The Economist 2019).

average welfare gains which, however, vary substantially across investors.Notably, older households with little financial wealth gain most, while younger households gain little due to human capital.

In more detail, identifying the effect of access to automated wealth management is challenging for two reasons. The first challenge is acquiring data. Regulatory filings, industry reports, and other public datasets do not contain information about the composition of households who participate with specific wealth managers, which is central to our paper. Accordingly, we obtain a novel dataset directly from a major U.S. robo advisor. We observe the demographic background, investment activity, and liquid assets (including retirement accounts), of households who participate with the advisor. This information enables us to empirically analyze the distributional effects of automated wealth management. In addition, we observe pairs of portfolios for households interested in becoming robo participants. Each pair contains the non-robo portfolio they manage themselves and the robo portfolio they would receive. We use this information to structurally decompose the channels through which automated wealth management improves welfare.

The second challenge is identification. We require a setting in which robo advice suddenly becomes more accessible to the modesty wealthy. We overcome this challenge by studying a quasi-experiment in which the same robo advisor unexpectedly reduces its account minimum from \$5,000 to \$500 in July 2015. This \$4,500 reduction constitutes a large shock for most U.S. households, as it equals 26% of the median U.S. household's liquid assets of \$17,000 at the time. The reduction enabled households with little investible wealth to access a suite of services typically reserved for the wealthy. We study the most basic of these services: a personalized, automatically rebalanced portfolio of risky assets. To the best of our knowledge, this shock is one of the first examples in which sophisticated wealth management becomes available to a wide range of non-affluent households.

We find that the reduction democratizes the market for automated wealth management. The wealth distribution of participants shifts sharply leftward after the reduction, while showing no pre-trend in the months leading up to it. In particular, the share of participants from the second and third U.S. wealth quintiles, whom we call the "middle class", increases by 107% (16 pps). This increase reflects a sharp break from trend that is not present among participants from the upper two quintiles, whom we call the "upper class". However, the democratization is asymmetric, in that there is no change in participation among the poorest quintile.

We formalize this graphical intuition through a difference-in-difference analysis. Our regres-

sion model compares the probability of participating with the robo advisor after versus before the reduction between middle versus upper-class households. Intuitively, the middle class represents the "treated" group in that it experiences a relaxation of minimum-account constraints due to the reduction. Accordingly, we find that middle-class households are 14 pps more likely to participate with the robo advisor after the reduction, relative to the upper class. We show that this estimate implies a 110% increase in the total number of middle-class robo participants. This finding is robust to various measures of wealth and, thus, is not biased by measurement error.

To check that we identify the desired effect, we examine whether a relaxation of minimum-account constraints drives the results. Otherwise, the middle class could have accessed wealth management without the reduction. Consistent with the constraints channel, the majority of middle-class households who became robo participants prior to the reduction bunched their investment at the previous minimum of $5,000, a hallmark of binding constraints. After the reduction such bunching immediately disappears, and most new middle-class participants make a previously infeasible investment of under $5,000. Additionally, based on a wide variety of tests, we find no evidence that the results are driven by channels distinct from a relaxation of constraints, such as: heightened visibility from targeted advertising or media attention (e.g., Kaniel and Parham (2017)); gambling motives (e.g., Bombardini and Trebbi (2012)); business stealing from competitors; or heterogeneous trends by demographics or risk attitude.

Our quasi-experimental results indicate significant middle-class demand for automated wealth management, but the channels that give rise to such demand remain unclear. As a natural step toward understanding these channels, we compare portfolios managed by households themselves versus those managed by the robo advisor. This comparison yields three insights. First, robo portfolios contain greater exposure to priced risk factors, such as bond and value premia. Consequently, they provide a 2 pps higher expected return than their self-managed counterparts. Second, robo portfolios contain strikingly less uncompensated risk. For example, based on the return-loss calculation of Calvet, Campbell and Sodini (2007), self-managed portfolios lose 7 pps in expected return from not investing with the same mean-variance efficiency as the robo advisor. Third, robo portfolios exhibit a "double glide path" by both age and wealth: a robo portfolio's systematic risk falls as a household grows older, holding wealth fixed; and its risk rises as a household become wealthier, holding age fixed. Therefore, robo portfolios are more personalized than target date funds (TDF) which offer a single glide path by age.

Motivated by these portfolio-level differences, we ask whether embedding them into a quan-

titative model of asset allocation can rationalize our empirical evidence. This exercise has two purposes. First, it tests whether middle-class households' demand for robo advisors reflects basic differences in asset allocation, or whether nonfinancial channels such as peace-of-mind are needed (e.g., Gennaioli et al. (2015)). Second, it allows for a more rigorous assessment of heterogeneous welfare gains. For example, the same gain in expected return may imply differential welfare gains for younger and older investors, after accounting for differences in labor income over remaining working life.

We add two realistic ingredients to benchmark models (e.g., Cocco, Gomes and Maenhout (2005)). First, we allow households to choose between managing their own portfolio or delegating it to a professional. Importantly, we do not make any prior assumptions about the differences between self-managed and delegated (i.e., robo) portfolios. Instead, we use our portfolio-level dataset to calibrate each portfolio's risk profile across the joint distribution of household age and wealth. The calibrated portfolio differences constitute a sufficient statistic, and so we do not need to take a stance on the behavioral friction generating these differences. The model's second distinguishing ingredient is an account minimum required by professional managers, which makes professional management less accessible. After adding these two ingredients, we ask whether the model can quantitatively explain our empirical results.

The model formalizes the following intuition. Suppose self-managed portfolios have lower expected returns and greater uncompensated risk than professionally managed portfolios, as in the data. Given these differences in portfolios, households optimally seek to allocate a share of their wealth to professional managers. However, middle-class households cannot achieve this optimal share because it requires an investment below the required minimum. Therefore, they can either invest more than their optimal share or simply not participate in wealth management. In the latter case, reducing the minimum relaxes minimum-account constraints and so prompts them to participate. The poorest households, however, may still find that participating in wealth management requires an excessive risky share. Consequently, the reduction has an asymmetric effect on participation across the wealth distribution, as we find empirically.

We show that the model quantitatively replicates our key empirical findings on both the extensive (e.g., growth in participation) and intensive margins (e.g., risky portfolio share). Thus, from the perspective of economic theory, we interpret the large empirical effect of the reduction as an optimal response to differences in portfolio characteristics. In particular, this result suggests that households act according to canonical financial motives (e.g., diversification, priced risk

exposure). While we do not rule out important nonfinancial motives (e.g., attention costs, peace-of-mind), these motives to not appear to be the dominant driver of household demand, after first conditioning on differences in asset allocation.

Turning to welfare, we focus on the distribution of welfare gains and the robo portfolio characteristics that generate these gains. Overall, the reduction raises welfare for new robo participants by 0.8%, based on the standard lifetime consumption metric. This increase is moderate. For reference, increasing permanent income by 5 pps with no reduction would similarly raise welfare by 0.9%. The moderate overall effect masks substantial heterogeneity. Across the age distribution, the reduction raises welfare three times as much for middle-class households over age 55 than for those under age 35. Intuitively, younger middle-class households have substantial human capital. Thus, despite their modest financial wealth, most of them will overcome the previous minimum within a few years. By contrast, older middle-class households have both modest financial wealth and modest human capital, and so their likelihood of eventually participating with the advisor is low.

To understand why households seek robo portfolios in the first place, we recalculate the welfare gain under a counterfactual in which robo portfolios do not possess certain characteristics. For example, robo portfolios are better diversified than self-managed ones, but, if they were not, then the reduction would result in a 70% (0.6 pps) lower welfare gain. Based on a similar calculation, removing robo portfolio's exposure to bond risk factors would also reduce the welfare gain by 70%. By contrast, differences in exposure to stock risk factors have a more modest impact on welfare, partly because households already have such exposure. Turning to dynamic asset allocation, robo portfolios resemble Target Date Funds (TDFs) in that they become less aggressive as the household ages. However, they also become less aggressive if the household's wealth declines. Both margins of personalization improve welfare, but, interestingly, personalization by wealth improves welfare as much as personalization by age. Lastly, incorporating a fixed cost of self-management a la Vissing-Jørgensen (2002) implies that the reduction partly raises welfare by bringing households into the stock market. This channel does not dominate, though, because stock market non-participation results in only a mild welfare loss relative to investing inefficiently on one's own.

Our results challenge several paradigms about the wealth management industry. First, wealth management is not necessarily a luxury good, and, in fact, middle-class households have substantial demand for it. Second, automated wealth managers provide the largest benefit to older

middle-class households, despite their popular image as an investment for the young (e.g., Hutchins (2020)). Third, robo portfolios do not simply substitute for a low-cost equity index fund, since much of the value from accessing robo portfolios comes from more sophisticated features like personalization and exposure to additional risk factors.

Lastly, we think carefully about how using data from a single robo advisor affects our conclusion. On one hand, this restriction offers a rare opportunity to study automated wealth management in its purest form, and, in that sense, it is ideal. Indeed, the advisor that we study has staunchly resisted the robo industry's trend toward a hybrid model that features a human manager. On the other hand, working with a self-selected sample limits inference about aggregate effects of automation in financial services. We evaluate this limitation by recalculating our main results after correcting for selection a la Heckman (1979). While this correction yields only mild differences, we nevertheless view any aggregate implications of our findings as suggestive (e.g., impact on wealth inequality).

We conclude this section by situating our contribution within the literature. Section 2 provides institutional background and describes our quasi-experiment. Section 3 describes our data. Section 4 estimates the effect of the reduction. Section 5 assesses robustness. Section 6 describes the model. Section 7 studies the model's descriptive implications. Section 8 studies welfare implications. Section 9 concludes.

Related Literature

Firstly, we contribute to a literature on FinTech intermediaries, and, specifically, to the nascent literature on robo advisors. We do so in three ways: institutional context, research question, and methodology. We contribute through our context by studying automated wealth management in a pure form that involves: complete portfolio delegation, as opposed to non-binding advice (e.g., D'Acunto, Prabhala and Rossi (2019); Bianchi and Briére (2020); D'Hondt et al. (2020)); no option for human interaction (e.g., Rossi and Utkus (2020)); and robo advice unaffiliated with the banking system (e.g., Loos et al. (2020)). We contribute through our research question by studying robo advisors' effect on the extensive margin. We can pursue this extensive margin analysis because our setting overcomes two hurdles in the literature: quasi-experimental variation in accessibility; and data on households' total liquid wealth.[4] Lastly, we contribute in methodology by

[4]Our extensive margin focus relates to a broader question within the FinTech literature of how technological innovation fosters financial inclusion or reduce wealth inequality. Most directly, our empirical results confirm the Philippon (2019) prediction that robo advising affects inequality by favoring the middle class over both the upper and lower

structurally analyzing the effects of robo advice. Applying this structural analysis to our quasi-experimental evidence enables us to rationalize households' demand for robo advice, to calculate the distribution of welfare gains, and to rank the channels that drive these gains.

Second, we contribute to the literature on fund managers by characterizing demand for passive management in a novel setting. Our focus parallels papers that reconcile the demand for active managers with their well-documented underperformance (e.g., French (2008)). For example, households may seek active managers because they may possess superior information (e.g., Gârleanu and Pedersen (2018)) or confer peace of mind (e.g., Gennaioli, Shleifer and Vishny (2015)). We find that households seek passive managers (i.e., robo advisors) largely because of an inability to take priced risk or diversify idiosyncratic risk on their own. In this sense, we corroborate the Von Gaudecker (2015) finding that Dutch households diversify better when they consult a financial professional. We also find that households value personalization in robo portfolios, complementing papers studying another form of personalized passive management, TDFs (e.g., Balduzzi and Reuter (2019); Parker, Schoar and Su (2021)). Specifically, households value how robo portfolios personalize according to age, as does a TDF, and, in equal measure, they value how robo portfolios personalize by wealth. Thus, robo portfolios may substitute for employer-sponsored TDFs when households lack access to them, which, as Parker et al. (2022) document, can significantly impact middle-class households' equity exposure.

Third, we contribute to the household finance literature empirically and through our model. Empirically, we characterize account minimums as a novel friction that limits investment in risky asset markets. This friction arises from the supply side and does not directly depend on household characteristics such as preferences (e.g., Barberis, Huang and Thaler (2006)), sophistication (e.g., Grinblatt, Keloharju and Linnainmaa (2011); Christelis, Jappelli and Padula (2010)), socialization (e.g., Hong, Kubik and Stein (2004)), or education (e.g., Cole, Paulson and Shastry (2014); Van Rooij, Lusardi and Alessie (2011)). Our empirical findings caveat Haliassos and Bertaut (1995) who, based on a two-period model, conclude that account minimums at discount brokers do not affect household investment. In terms of model, we follow in a long tradition of quantitative portfolio choice models summarized by Gomes (2020). Relative to these models, we match quasi-experimental evidence, incorporate both self-managed and professionally managed portfolios, and, like Fagereng, Gottlieb and Guiso (2017), calibrate it using microdata.

classes. More broadly, see analogous results on financial inclusion in the contexts of app-based payments (e.g., Hong, Lu and Pan (2020)), bank deposits (e.g., Bachas et al. (2018); Bachas et al. (2020); Higgins (2020)), and mortgage markets (e.g., Fuster et al. (2019); Bartlett et al. (2021); Fuster et al. (2021)).

2 Institutional Background

This section describes the U.S. robo advising market (2.1), the advisor we study (2.2), and our quasi-experiment (2.3). To clarify our terminology, we use "robo advisor" and "automated wealth manager" synonymously, even though the latter is nested within the former. In particular, the robo advisor we study is an automated wealth manager because it offers an entire pipeline of wealth management services ("wealth manager") without human interaction ("automated"). This includes identifying investment goals, building a personalized risk profile, choosing the appropriate portfolio allocation, managing and rebalancing this portfolio on a periodic basis, and other services that we do not study in this paper.

2.1 The U.S. Robo Advising Market

Paraphrasing D'Acunto and Rossi (2020), robo advisors emerged in the mid-2000s in response to the limitations of traditional wealth managers. They are distinguished by relying on algorithms to select and maintain an allocation for their clients. This automated approach features lower per-portfolio management costs relative to the traditional approach of manually constructing and managing a client's portfolio (Moulliet et al. (2016)). In practice, several robo advisors also incorporate human judgment on a portfolio-by-portfolio basis, much as a traditional manager would. Others rely purely on algorithm, including our data provider, Wealthfront.

Appendix Table A1 summarizes the largest robo advisors in the U.S. as of July 2015, including their account minimums, assets under management, fees, and provision of traditional, human-based management. At the time of our analysis, Wealthfront managed roughly $3 billion and was the largest standalone robo advisor in the U.S. market, with Betterment and Personal Capital as its nearest competitors. Two traditional managers, Vanguard and Charles Schwab, launched robo advising services early in 2015. Both of these services managed more than Wealthfront because they transferred assets from existing, non-robo services.

Notably, Wealthfront is the only robo advisor that relies purely on automation, with no option for a human advisor. Indeed, throughout its history, Wealthfront has marketed itself as a technology-oriented firm, and it has staunchly opposed the industry's trend toward a hybrid robo advising model that does not rely strictly on automation. Thus, from the standpoint of learning about automated wealth management in its purest form, using data from a single robo advisor actually makes our analysis more representative.

2.2 The Robo Advisor

Wealthfront, henceforth "the robo advisor", has offered many services throughout its history, including tax loss harvesting, long term financial planning, portfolio lines of credit, and a risk parity fund. Most relevant for this paper is its baseline product: an automatically rebalanced portfolio of 10 ETFs corresponding to 10 asset classes.[5] The portfolio weights are determined by a questionnaire that asks the client several questions about age, liquid assets, income, demographic background, and response to hypothetical investment decisions. The client is then assigned to one of 20 possible risk tolerance scores, which range from 0.5 to 10 in increments of 0.5. Each risk tolerance score uniquely determines a robo portfolio. The portfolio weights solve a problem of optimal asset allocation across the 10 ETFs, taking this score as a parameter.

Figure 4 examines the risk profile of robo portfolios. Overall, this risk profile conforms to "textbook" recommendations for retail investors (e.g., Malkiel 2015). First, robo portfolios with a higher risk tolerance score exhibit greater total volatility, shown by the solid line in panel (a). After separating idiosyncratic risk from priced risk using a standard five-factor pricing model, we find that most of the volatility in robo portfolios reflects diversified exposure to priced risk factors, shown by the dashed line.[6] Panel (b) examines how priced risk varies by age and wealth. Robo portfolios reduce their risk as households age. Moreover, for a given age, robo portfolios provide less risk for less-wealthy households than for wealthier ones. These patterns may reflect the robo advisor's desire to accommodate life cycle effects and precautionary savings motives, respectively. The robo advisor implements these glide paths by reducing stock exposure for older and less-wealthy households and, symmetrically, increasing their bond exposure, shown in panels (c) and (d). Lastly, Appendix Figure D shows how robo portfolios reduce their exposure to value stocks as households age, consistent with the Jurek and Viceira (2011) solution.

Importantly, robo portfolios are not recommendations: rather, they are directly managed by the robo advisor. Consequently, households have little discretion over their portfolio allocations, and so their robo performance will not depend on sophistication (e.g., Grinblatt, Keloharju and Linnainmaa (2011); Christelis, Jappelli and Padula (2010)), ability to diversify (e.g., Calvet, Camp-

[5]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. The 10 primary ETFs are chosen to track stock market indices (VIG, VTI, VEA, VW), bond market indices (LQD, EMB, MUB, TIPS), and other asset classes, namely real estate (VNQ) and commodities (XLE).

[6]The factors are the three Fama and French (1993) factors plus U.S. and global excess bond returns. Appendix D contains details. Idiosyncratic volatility and volatility from factor exposure do not sum to total volatility by Jensen's inequality.

bell and Sodini (2007)), willingness to follow advice (e.g., Bhattacharya et al. (2012)), or reluctance to rebalance (e.g., Calvet, Campbell and Sodini (2009)).

2.3 The 2015 Reduction in Account Minimum

On July 7, 2015, the robo advisor unexpectedly reduced its account minimum from $5,000 to $500, meaning that a household would need to invest only $500 to participate with the advisor as opposed to $5,000 beforehand. This reduction is quite sizeable from the standpoint of most U.S. households. For reference, $5,000 equals 30% of the median household's liquid assets ($17,000), and it defines the 37th percentile of the U.S. wealth distribution, according to the 2016 Survey of Consumer Finances. Prior to the reduction, therefore, half of U.S. households could not participate without investing at least 30% of their wealth, while 37% could not participate at all without borrowing. The reduction was motivated by the advisor's philosophy of inclusive investment and belief that non-wealthy households will eventually accumulate enough assets to become high-revenue customers.[7] Indeed, given the advisor's management fee of 0% for accounts under $10,000 (0.25% for larger accounts), the reduction was not intended to raise short-term revenue.

At the time of the reduction, all of the largest five U.S. robo advisors required an account minimum of at least $5,000 except for one, which had no account minimum but maintained a fee structure that discouraged setting up small accounts.[8] Importantly, the month of the reduction does not coincide with any other product launches by the robo advisor, any changes in its fee, or any significant developments in the overall robo advising market. This effectively idiosyncratic timing allows us to identify the reduction's effect on household participation in automated wealth management in Section 4.

We interpret the reduction as a shock that expands access to automated wealth management, rather than as a direct effect of automation itself. That said, automation quite plausibly enabled the reduction by reducing fixed costs of portfolio management. For example, a single manager can oversee 330 times as many automated portfolios as non-automated ones (Moulliet et al. (2016)). Thus, in asking whether automation affects inequality in wealth management, we principally mean whether expanded access to automated wealth management affects such inequality, using

[7] In the words of the robo advisor's then-CEO: "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."

[8] Betterment charged a $3 service fee on accounts under $10,000 for customers who do not auto-invest $100 monthly in their accounts. This fee structure implies a 7.2% annual management fee for a $500 account and a 36% management fee for a $100 account (Thomson Reuters 2015).

the reduction as a source of variation in "expanded access". We leave open the likely possibility that automation actually enabled the reduction itself.

3 Data

Our core analysis relies on two datasets that draw from the same population: a panel dataset covering deposit activity by households who participate with the robo advisor (3.1); and a dataset on self-managed, non-robo portfolio holdings that includes both households in the deposits dataset and robo non-participants (3.2). We now describe these two datasets, other auxiliary datasets (3.3), and summary statistics (3.4). Appendix A has details. For the rest of the paper, we use the term "robo participant" to describe households who have invested money with the robo advisor.

3.1 Deposits Dataset

The first core dataset contains a weekly time series of deposits with the robo advisor from December 1, 2014 through February 29, 2016. This window straddles the July 2015 reduction in minimum. It also marks a formative period in the history of the robo advising market when the number of participating households was still small. We obtained this dataset through a direct query of the robo advisor's internal server. Thus, we observe the same information as would an analyst working for the advisor. Specifically, we observe the date and size of the deposit, whether the deposit comes from a new participant with the robo advisor, and the following demographic variables about the participating household: annual income; state of residence; household age; and liquid assets, defined as "cash, savings accounts, certificates of deposit, mutual funds, IRAs, 401ks, and public stocks". The demographic variables are self-reported via the robo advisor's questionnaire and static. Thus, liquid assets may be subject to measurement error from misreporting, but we check in Section 4.3 that such measurement error does not bias our results.

Studying a company-specific dataset has two advantages over publicly available datasets such as the SEC's Form ADV filings, which serve as the basis for many industry reports about the robo market. First, estimates of robo participation growth derived from Form ADV data would be highly imprecise because they can include both inactive clients and "clients" who create a username but never provide the robo advisor with any money.[9] Second, unlike public data, our

[9]For example, we observe 9,702 participants in our dataset, in contrast to the 61,000 reported in publicly available SEC filings. This discrepancy reflects how: "The definition of 'client' for Form ADV states that advisors must count clients who do not compensate the advisor" (SEC 2017).

dataset includes information about a robo participant's wealth, allowing us to study investment activity across the wealth distribution.

3.2 Portfolio Dataset

The second core dataset contains snapshots of self-managed, non-robo portfolio holdings for a subset of households in the deposits dataset and for a set of households who consider participating with the robo advisor but choose not to do so. The advisor obtains these snapshots from an online tool through which it provides free financial advice to candidate clients about their outside portfolio holdings. As with the deposits dataset, we obtained this portfolio dataset through a direct query of the robo advisor's internal server. We merge this dataset with security-level information from standard sources (e.g., CRSP) to produce a cross-section of 1,913 portfolio pairs. Each pair consists of a self-managed portfolio and a counterfactual robo portfolio that the candidate client would receive by becoming a robo participant.

We use the portfolio dataset to calibrate the model in Section 6, and three features make the dataset ideal for this purpose. First, we observe overall advisory and management fees. By restricting the set of non-robo portfolios to those without such fees, we can ensure with a high degree of confidence that the non-robo portfolios we study are managed by their owner (i.e., self-managed). Using the portfolio dataset, we can realistically parameterize the choice between managing one's own portfolio versus delegating it to a professional. Second, we observe the self-managed portfolios not only of robo participants (45% of sample), but also of robo non-participants. Observing non-participants helps our calibration avoid selection bias from, say, the possibility that only households who cannot invest efficiently on their own delegate to a professional. Third, we observe portfolios which non-participants would receive if they too decided to become participants. Chalmers and Reuter (2020) highlight the importance of observing counterfactual portfolios when evaluating demand for asset allocation and security selection.

Summarizing, each dataset accomplishes a separate goal, but both draw from the same population. The deposits dataset enables our main empirical analysis: to examine how expanded access to automated wealth management affects the wealth distribution of robo participants. The portfolio dataset enables us to calibrate the model that explains our empirical results, since it includes information on robo and non-robo portfolios for both robo participants and non-participants. In terms of external validity, the underlying population for both datasets consists of households who endogenously know about the robo advisor. This sampling restriction does not limit our ability to

learn about the robo market since, for that purpose, these households are exactly those we seek to understand. However, for the broader purpose of informing public policy, this restriction would require us to correct for how such households do not represent the broader U.S. population. We do so in Section 5.4, but we are nevertheless cautious in our policy statements.

3.3 Auxiliary Datasets

The most important auxiliary dataset is the 2016 Survey of Consumer Finances (SCF). This dataset includes financial and demographic information about a representative cross-section of U.S. households, as summarized by Bricker et al. (2017). The SCF enables us to benchmark a household's wealth in our robo advising dataset against the U.S. population. We respectively use the terms "lower class", "middle class", and "upper class" to describe households from the first, second or third, and fourth or fifth quintiles of the overall U.S. distribution of liquid assets, where liquid assets are calculated to match the definition in our robo advising dataset as closely as possible. The precise definition is in Appendix A.3. The resulting boundary between the lower versus middle class is $1,000 in liquid assets, and the boundary between the middle versus upper class is $42,000 (Appendix Table A3).

3.4 Summary Statistics

Table 1 summarizes the deposits dataset. Panel (a) compares households who become robo participants after the reduction in account minimum (i.e., new participants) with existing participants. New participants are significantly less wealthy, earn lower incomes, make smaller initial deposits, and are 16 pps more likely to belong to the middle class, relative to existing participants. They are also more likely to come from less financially developed U.S. states, as shown in Appendix Figure A1.

Panel (b) conveys a similar pattern when restricting the comparison to middle-class households. The median new middle-class participant's initial deposit of $2,000 would have been infeasible under the previous account minimum of $5,000. Indeed, over half of existing middle-class participants invested exactly $5,000 for their initial deposit. This "bunching" suggests that the previous minimum imposed a binding constraint on their investment behavior. Importantly, new middle-class households invest with persistence. For example, 97% do not close their account over our sample period, and, in fact, 72% make a subsequent deposit. Such subsequent deposit-making

resembles the "dollar cost averaging" strategy commonly advocated by practitioners, which Brennan, Li and Torous (2005) show is optimal for risk averse investors.

Lastly, new middle-class participants are not significantly younger than existing ones. This suggests that the reduction does not work through generation-specific effects, such as technological savviness. Instead, it could be that older middle-class households would like to invest with the robo advisor, just like younger middle-class households, but the previous minimum had constrained their ability to do so. Section 8.1 substantiates this conjecture using standard life cycle logic. In particular, older middle-class households value how robo portfolios reduce their risk both with age and after a decline in wealth, as already documented in Figure 4.

4 Democratization of the Robo Market

We examine how the reduction in minimum affects robo participation by constrained, middle-class households. First, we provide graphical evidence (4.1). Then, we formalize our identification strategy (4.2). We report our main results based on a first-difference estimator (4.3) and a dynamic difference-in-difference estimator (4.4).

4.1 Graphical Evidence

Two pieces of graphical evidence suggest that reduction democratizes the robo market by relaxing minimum-account constraints on the middle class. First, Figure 1 plots the distribution of new robo participants by time period. This distribution shifts sharply leftward in the quarters after the reduction. Moreover, it remains stable in the preceding quarters, providing preliminary evidence that the post-reduction shift does not simply reflect a pre-trend.

Second, Figure 2 shows how the leftward shift in Figure 1 makes the robo wealth distribution more representative of the overall U.S. wealth distribution, "democratizing" the robo market. Specifically, the share of robo participants from the second and third quintiles of the U.S. wealth distribution grows by 107% (16 pps), while the share from the upper two quintiles falls by 18% (16 pps). However, there is a non-monotonic relationship between robo participation growth and wealth, since households from the first quintile remain non-participants. Appendix Figure A4 replicates this finding using the subset of robo participants who are in the portfolio dataset.

4.2 Identification

Begin with the following flexible model of robo participation in period $\mathcal{T}$,

$$Participant_{i,\mathcal{T}} = \mu\left(Middle_i \times Post_{\mathcal{T}}\right) + \psi\left(X_i \times Post_{\mathcal{T}}\right) + \zeta_i + \varrho Post_{\mathcal{T}} + v_{i,\mathcal{T}}, \tag{1}$$

where i indexes household; $\mathcal{T}$ indexes the pre-reduction period (i.e., $\mathcal{T} = 0$) versus the post-reduction period (i.e., $\mathcal{T} = 1$); $Participant_{i,\mathcal{T}}$ indicates if household i participates with the robo advisor at some point in period $\mathcal{T}$; $Middle_i$ indicates if i belongs to the second or third U.S. wealth quintile, in contrast to the fourth or fifth quintiles that comprise the reference group; ζ_i is a household fixed effect; and X_i is a vector of household characteristics: age, log income, state of residence fixed effects, and a measure of subjective risk tolerance from the advisor's questionnaire.

We propose that the reduction affects robo participation among households with moderate levels of wealth because it relaxes constraints on their ability to invest. Equation (1) measures "moderate wealth" using the indicator $Middle_i$. Therefore, under an identification assumption described shortly, the parameter μ equals the effect of the reduction on middle-class households' probability of robo participation. Explicitly, μ equals the double difference in the probability of becoming a robo participant after versus before the reduction between middle-class versus upper-class households.

The additional terms in equation (1) capture channels distinct from the minimum-account constraints channel. The fixed effect ζ_i captures slow-moving characteristics that predispose households to participating with the advisor, such as sophistication or trust. Since such "affinity" to the advisor increases the probability of participation in any period, we can separately identify the effect of minimum-account constraints because the minimum changes over time. The interaction between X_i and $Post_{\mathcal{T}}$ captures heterogeneous trends by observed household characteristics. If, for example, younger households are more likely to become robo participants after the reduction for reasons apart from a relaxation of minimum-account constraints, then the coefficient ψ would separately capture this effect.

Estimating equation (1) is equivalent to estimating the first-differenced equation,

$$\Delta Participant_i \equiv New\ Participant_i = \mu Middle_i + \psi X_i + \varrho + u_i, \tag{2}$$

where $New\ Participant_i$ indicates if household i becomes a robo participant after the reduction;

and $u_i \equiv \Delta v_i$. We estimate equation (2) on the set of eventual robo participants in the deposits dataset. Therefore, μ equals the reduction's effect on the probability of robo participation, conditional on eventually participating.[10] In an extension, we also estimate equation (2) using the portfolio dataset, in which we need not condition the interpretation of μ on eventual participation because we observe non-participants as well.

The following identification assumption allows us to interpret μ as the causal effect of the reduction on middle-class households' probability of robo participation:

$$0 = \mathbb{E}\left[Middle_i \times u_i \middle| X_i \right]. \tag{3}$$

Equation (3) states that unobserved determinants of a *change* in robo participation, u_i, do not systematically vary across the middle and upper classes, conditional on the household's observable characteristics, X_i. This implies that the difference in the change in robo participation between the middle and upper classes reflects the effect of a lower account minimum.

There are three ways in which equation (3) could be violated. The first is measurement error in self-reported liquid assets, discussed at length below. The second source of bias are omitted variables that are not caused by the reduction but nevertheless affect changes in middle-class households' robo participation. One such confounding variable could be trend growth in households' awareness of robo advisors. Such bias is unlikely, given the strong parallel trends documented shortly in Figure 3 and Appendix Figure D. Another potentially confounding factor could be contemporaneous developments in the robo industry, such as the launching of new robo products by the two traditional managers named in Section 2.2. However, these new products were not targeted toward the middle class, and they were launched at least two months prior to the reduction, comfortably before the strong divergence in middle-class households' behavior. The third source of bias are omitted variables that are themselves caused by the reduction. The leading examples are targeted media attention or advertising. If middle-class households are more exposed to such media or advertising, then μ confounds the effect of heightened visibility with the effect of minimum-account constraints. In Section 4.4, we test for bias from heightened visibility and find

[10] One can also interpret μ as the share of robo participants whom the reduction causes to participate. Clarifying this point, the reduction causes a household to participate if she finds it optimal to participate under the \$500 minimum but not under the \$5,000 minimum. Let $Constrained_i$ indicate such a household. The share of eventual participants whom the reduction causes to participate is thus: $\Pr[Constrained_i = 1 \mid Participant_{i,1} = 1]$. By definition, the reduction cannot cause an existing participant to participate. Morover, by assumption (3), the reduction does not affect upper-class households. Therefore, the previous two statements imply: $\Pr[Constrained_i = 1 \mid Participant_{i,1} = 1] = \Pr[New_i = 1 \mid Participant_{i,1} = 1] - \Pr[New_i = 1, Middle_i = 0 \mid Participant_{i,1} = 1] = \mu$, where the last equality follows from equation (2) after omitting controls X_i for brevity.

that its effect is very modest and does not interfere with our results.

4.3 First-Difference Results

Table 2 reports the results. The estimate in column (1) implies that middle-class households are 15 pps more likely to become robo participants after the reduction in account minimum than before the reduction, relative to upper-class households. Column (3) shows that the effect of the reduction does not vary with age, in line with the descriptive evidence in Table 1. We later interpret this finding from the perspective of a portfolio choice model in Section 8. In column (4) we estimate equation (2) using the portfolio dataset, which includes non-participants. The result implies that the reduction raises middle-class households' probability of participation by 8.6%, without conditioning on eventual participation. The smaller effect may reflect how middle-class households in the portfolio dataset may be relatively sophisticated and, thus, have less demand for outside management.

Our treatment exposure variable, $Middle_i$, may be subject to additive measurement error due to self-reporting. As we show formally in Appendix B, such measurement error tends to bias the estimate toward zero (i.e., attenuation bias). The exception is if new participants underreport their wealth relative to existing participants. We mitigate this concern by remeasuring $Middle_i$ in two ways. First, we redefine the middle class exclusively as the second quintile of the U.S. wealth distribution and omit households from the third quintile from the sample. Under this definition, upper-class households would need to underreport liquid assets by at least \$36,000 to be misclassified as middle-class. Second, we exclude households whose liquid assets are within a 10% buffer of the boundary between the third and fourth quintiles. This approach removes all cases of mismeasurement less than \$8,400 ($2 \times 0.1 \times 42,000$). Columns (5) and (6) show that the estimates based on these alternative measures of $Middle_i$ equal 15%, the same as our baseline estimate. Thus, the baseline estimate does not seem to be biased by measurement error.

In our main tests, standard errors are clustered by household (i.e. heteroscedasticity-robust). Appendix Table A4 shows how the main results remain statistically significant when clustering standard errors by state of residence.

4.4 Dynamic Difference-in-Difference Results

The first-difference equation (2) has the advantage of parsimony, but it does not inform whether

higher-frequency dynamic effects influence the baseline results (e.g., heterogeneous pre-trends, media attention, or advertising). Accordingly, we estimate the dynamic difference-in-difference (DiD) equation

$$New\ Participant_{i,t} = \mu\left(Middle_i \times Post_t\right) + \zeta_i + \varrho_t + \psi\left(X_i \times Post_t\right) + u_{i,t}, \tag{4}$$

where i and t index household and week; $Post_t$ indicates if t is greater than the week of the reduction; $New\ Participant_{i,t}$ indicates if i becomes a robo participant in week t, as opposed to the other weeks in our observation window; ζ_i is a household fixed effect; and ϱ_t is a week fixed effect. The standard errors are clustered by household and week.

Since the outcome variable in equation (4) differs from its counterpart in equation (2), the interpretation of the parameter μ also changes. In particular, μ now equals the effect of the reduction on middle-class households' probability of robo participation in any given week rather than the cumulative effect over the 32-week period after the reduction. The estimates in columns (7)-(8) of Table 2 lie between 0.6 pps and 0.8 pps. These estimates imply the cumulative port-reduction effects between 19.2 pps (0.6×32) and 25.6 pps (0.8×32) which are slightly larger than the first-difference results.

We further rigorously evaluate dynamic confounds through an event study that builds off of equation (4). Explicitly, we interact $Middle_i$ in equation (4) with a vector of month fixed effects to trace out the reduction's full dynamic effect, estimating the following specification:

$$New\ Participant_{i,t} = \sum_{m \neq June\ 2015} \left(\mu_m \times Middle_i \times \mathbb{1}_{t \in m}\right) + \zeta_i + \varrho_t + \sum_{m \neq June\ 2015} \left(\psi_m \times X_i \times \mathbb{1}_{t \in m}\right) + u_{i,t}, \tag{5}$$

where μ_m's are the coefficients on the treatment indicator, $Middle_i$, that vary non-parametrically by month m which includes week t. We omit the month prior to the reduction (June 2015) such that the μ_m's can be interpreted relative to this baseline period. In particular, these coefficients represent the difference in outcomes between the middle-class and upper-class households in each month relative to June 2015. The parallel trend assumption implies thats there is no difference between the households prior to the reduction (i.e. $\mu_m = 0$ for all months prior to July 2015).

The results in Figure 3 show that the probability of becoming a robo participant increases

sharply and significantly for middle-class households exactly in the month of the reduction, July 2015, consistent with Figure 1. Additionally, the differences between the middle-class and upper-class households are economically small and statistically indistinguishable from zero before the reduction, providing strong support for the parallel trend assumption. It is thus unlikely that pre-trends in middle-class households' robo participation or other dynamic effects bias the baseline results. Otherwise, such confounding factors would need to coincide exactly with the month of the reduction, which, per the institutional background in Section 2, is highly unlikely.

5 Robustness

We assess the internal and external validity of the reduced-form results. Specifically, we directly assess the minimum-account constraints channel (5.1), evaluate media attention and advertising (4.4), perform an aggregation exercise to assess magnitude (5.3), assess external validity (5.4), and discuss various other potential forms of bias, such as business stealing or gambling motives (5.5). The results of all these tests support the baseline results' validity.

5.1 Testing the Constraints Channel

An implicit corollary of assumption (3) is that the reduction affects middle-class households' participation by relaxing constraints on their investment with the robo advisor. Table 1 already provides two summary statistics consistent with the constraints channel: 55% of existing middle-class robo participants bunched their initial deposit at the previous minimum ($5,000); and 66% of new robo participants make an initial deposit under the previous minimum.

We formalize this intuition through a DiD estimator and an event study. First, we estimate a variant of the DiD equation (4) with a different outcome variable that measures bunching around the previous minimum. Explicitly, we estimate

$$\mathbb{1}[\textit{First Deposit}_{i,t} = \$5k] = \mu\,(\textit{Middle}_i \times \textit{Post}_t) + \varrho_t + \psi_t X_i + u_{i,t}, \tag{6}$$

where $\mathbb{1}[\textit{First Deposit}_{i,t} = \$5k]$ indicates whether the initial deposit made by household i in week t equals the previous account minimum or is no more than 5% higher; and the sample consists of household-weeks in which the household makes its initial deposit, so that we cannot include household fixed effects.

Table 3 reports the results. Focusing on column (3), the estimated coefficient on $Middle_i$ implies that middle-class households who become participants before the reduction are 20 pps more likely to invest right at the minimum, relative to upper-class households. Such bunching is a hallmark of constrained behavior. Consistent with a subsequent relaxation of constraints, this propensity falls by 24 pps after the reduction, per the estimated coefficient on $Middle_i \times Post_t$. The point estimate is roughly 50% larger for households from the second wealth quintile ($Lower\ Middle_i$) than the third quintile ($Upper\ Middle_i$). Since minimum-account constraints bind more strongly for less-wealthy households, this heterogeneity further supports the constraints channel. Collectively, these results suggest that middle-class households experience a relaxation of constraints, thereby supporting assumption (3).

5.2 Media Attention and Targeted Advertising

Media attention, advertising, or other changes in visibility around the reduction may disproportionately influence middle-class households, thus biasing our baseline results upward. The ideal approach to this problem would incorporate data on advertisements from all media sources observed by each household i. In the absence of such an ideal advertising dataset, we use several time-varying proxies, denoted Adv_t. Then, we re-estimate the DiD equation (4) after including the interaction $Middle_i \times Adv_t$. This interaction allows middle-class households to respond differently to the dynamic visibility of the robo advisor. Furthermore, we allow for this responsiveness to increase after the reduction (e.g., targeted advertising) by also including the triple interaction $Middle_i \times Post_t \times Adv_t$.

We define Adv_t in two ways, which require us to collect additional data on articles from Google News and on blog posts written by the robo advisor itself. The first definition is the number of news articles about the advisor published in the month of week t. The second definition is the number of blog posts written by the advisor in the month of week t. It is reasonable to assume that these proxy variables are positively correlated with unobserved advertisements across many other media sources. For example, when the company launches a new ad campaign, it simultaneously creates a series of blog posts on its website, promotes news articles, and runs ads across social media and other platforms.

Table 4 reports the results. In columns (1) and (3), we estimate a coefficient of around 0.8 on $Middle_i \times Post_t$, similar to the estimate in Table 2. Note that this coefficient is now interpreted as the effect of the reduction on the middle class when the level of advertising equals zero. The coefficient

on the triple interaction is interpreted as the marginal effect of advertising on the middle class after the reduction.

We next account for advertising that varies by the household's demographics (e.g. income) by interacting $Post_t$, Adv_t, and $Post_t \times Adv_t$ with the controls in Table 2. For the measure based on Google News the estimated coefficient equals 0.5 (column (2)), which is 0.1 pps lower than the estimate in the analogous specification in column (7) of Table 2. This difference suggests modest upward bias from targeted advertising. For the measure derived from the advisor's blog posts, we do not find any difference relative to the baseline specification (column (4)). Lastly, we account for visibility that varies geographically by including a similar set of interactions with a vector of fixed effects for state of residence. The results in columns (3) and (6) are no different.

While the effect of targeted advertising may modestly affect our estimates, we find no evidence of any major bias that may change our baseline conclusions.For example, we can calculate the cumulative post-reduction effect relying on the most conservative estimate after accounting for the advertising. Our smallest estimate equals 0.48 pps (column (3) of Table 4). In this case, the effect of the reduction on the probability of robo participation among the middle class over the 32 post-reduction weeks equals 15.36 pps (0.48×32), which is highly consistent with our baseline results from Table 2.

5.3 Interpreting the Magnitude of Effect

To sharpen the interpretation of the estimates in Table 2, we use them to back out the reduction's effect on the overall number of robo participants. Appendix B describes this procedure.[11] The result shown in Appendix Table A5 implies that the reduction increases the overall number of robo participants by 13%. This is driven by a 108% increase in the number of middle-class participants. Thus, the reduction seems to have spurred significant growth in overall participation with the robo advisor. Our paper focuses on the reduction's distributional effects as opposed to its aggregate effects, and so we do not attempt to aggregate these results to the market level. However, we note that most of the robo advisor's competitors have subsequently reduced their own account minimums. Therefore, it is plausible that increased participation by the middle-class has played an important role in the robo market's rapid growth.

[11] The reduction increases the number of participants by $\eta \equiv g - \frac{\psi \mathbb{E}[X_i]+\varrho}{1-(\psi \mathbb{E}[X_i]+\varrho)}$, where g is the observed growth rate in the number of robo participants from the pre-reduction period to the post-reduction period. The remaining notation is the same as in equation (2).

5.4 External Validity of Quasi-Experiment

By construction, the households in our data know about the robo advisor. We ask whether performing a similar quasi-experiment (i.e., reduction in minimum) among a broader population would yield a similar estimate of μ. Explicitly, we would like to reestimate equation (1) on a sample of eventual robo participants that includes both participants who endogenously know about the advisor, as in our data, and participants who do not, as might be found in a clinical trial that recruits from the U.S. population. Short of this ideal exercise, we instead model the endogenous knowledge of the robo advisor explicitly and adjust our estimates similarly to Heckman (1979). We describe this procedure in Appendix B.3. The resulting first-difference estimate of μ falls by 3 pps to 12.5% after accounting for sample selection. This decline implies that households who know about the robo advisor are more likely to respond to the reduction. At the same time, the magnitude of the decline is small.

5.5 Business Stealing and Gambling

In Appendix B, we find no evidence of bias from a reallocation of participants across robo advisors (i.e., business stealing) or gambling motivations.

6 Description of Model

Introducing a model addresses two questions our reduced-form analysis cannot address alone. First, we would like to know whether middle-class households' estimated demand for wealth management can be rationalized by the differences in portfolios shown in Figure 4. As a natural starting point, we embed our empirical setting in a benchmark model of portfolio choice. Second, if our empirical results are indeed drive by portfolio differences, then the average middle-class robo participant must experience a welfare gain. We would like to examine the characteristics of robo portfolios that generate this gain, as well as the distribution of welfare gains across households. In this section, we describe the model's setup (6.1) and calibration (6.2). Results are in Sections 7 and 8.

6.1 Setup

We follow the structure of benchmark portfolio choice models as closely as possible (e.g.,

Campbell et al. (2001); Cocco, Gomes and Maenhout (2005)), with two principal additions. First, rather than investing in a single, perfectly diversified risky asset, households have two investment opportunities: a self-managed portfolio ($\mathcal{S}$); and a portfolio overseen by a wealth manager ($\mathcal{A}$). A priori, we take no stance on the characteristics of these portfolios. We instead let the data inform these characteristics, as described in Section 6.2. Second, the two portfolios differ in that the latter requires an account minimum, M. We soon narrow our focus to the particular, automated wealth manager described in Section 2. Therefore, one can imagine there is a broader set of unmodeled portfolios overseen by wealth managers, and portfolio $\mathcal{A}$ is the one with the lowest account minimum, that is, the robo portfolio.

6.1.1 Preferences

As in our empirical analysis, let i index household. Time is discrete, and t indexes year. For the rest of the exposition, we conserve notation by aligning a household's age with the year, such that we do not maintain both age and time subscripts. Households begin their problem at age t_0. With probability p_t, a household of age t survives until age $t+1$, and at age $\overline{T}$ any surviving households leave the model. Households consume $C_{i,t}$ each year. They have isoelastic preferences over flow consumption, with coefficient of relative risk aversion γ. Thus, household i of age t has expected lifetime utility

$$U_{i,t} = \mathbb{E}_t \left[\sum_{\tau=t}^{\tau=\overline{T}} \delta^{\tau-t} \left(\prod_{j=t}^{j=\tau-1} p_j \right) \frac{{C_{i,\tau}}^{1-\gamma}}{1-\gamma} \right], \tag{7}$$

where δ is the discount factor. The absence of a bequest motive in equation (7) improves the model's parsimony, since we do not need such a motive to match the data. Households enter age t with consumable resources $W_{i,t}$, frequently called "cash-on-hand" in the literature (e.g., Deaton (1991)). Cash-on-hand is replenished through income from financial assets and labor income, both of which we describe below. To match our empirical work, we call $W_{i,t}$ "liquid assets", since it governs not only how much a household can consume, but also how much she can invest. Our setup to this point falls very much in line with benchmark models.

6.1.2 Financial Assets

There are three financial assets: a risk-free asset, which gives return R^f and can be likened to a savings account; a risky self-managed portfolio, which gives return $R^{\mathcal{S}}_{i,t}$ and can be likened to

a discount brokerage account; and a risky portfolio overseen by an automated wealth manager, which gives return $R^{\mathcal{A}}_{i,t}$. The last of these portfolios requires an account minimum of M and a small management fee equal to that described in Section 2.2. We simplify the model's computational complexity by assuming households cannot hold the self-managed and automated portfolios concurrently. This simplification has little bearing on the results because it is rarely optimal for households to hold both portfolios at the same time.

We introduce a factor structure for risky returns. This approach will improve the quality of our calibration, as it addresses the well-known challenge of estimating expected returns in finite samples (e.g., Merton (1980)). Explicitly, we suppose the return on portfolio $\mathcal{P} \in \{\mathcal{S}, \mathcal{A}\}$ evolves according to

$$R^{\mathcal{P}}_{i,t} = \beta^{\mathcal{P}}_i F_t + \epsilon^{\mathcal{P}}_{i,t}, \tag{8}$$

where F_t is a vector of priced risk factors, normally distributed with mean π^F and covariance matrix Σ^F; $\beta^{\mathcal{P}}_i$ is the loading of portfolio $\mathcal{P}$ on F_t for household i; and $\epsilon^{\mathcal{P}}_{i,t}$ is an idiosyncratic shock, normally distributed with mean zero and volatility of $\sigma^{\mathcal{P}}_{\epsilon,i}$. The quantity of compensated risk, $\beta^{\mathcal{P}}_i$, may vary not only across portfolios, per the superscript $\mathcal{P}$, but also across households, per the subscript i. This flexibility can capture how, for example, robo portfolios become less risky as households age. Likewise, the quantity of uncompensated risk, $\sigma^{\mathcal{P}}_{\epsilon,i}$, may vary across both portfolios and investors.

6.1.3 Labor Income

Households retire at age $\underline{T}$+1. For $t \leq \underline{T}$, they receive uninsurable labor income, $Y_{i,t}$. Following the literature's convention (e.g., Carroll (1997)), labor income in years without a disaster evolves according to

$$\log(Y_{i,t}) = f_i + \xi_{i,t} + \nu_{i,t}, \tag{9}$$

where f_i is a deterministic function of age; $\xi_{i,t}$ is a transitory shock, normally distributed with mean zero and volatility of σ_ξ; and $\nu_{i,t}$ is a permanent shock that evolves according to

$$\nu_{i,t} = \nu_{i,t-1} + \underbrace{\Xi_t + \omega_{i,t}}_{\upsilon_{i,t}}, \tag{10}$$

where $v_{i,t}$ is normally distributed with mean zero and volatility of σ_v. Equation (10) implies that permanent income shocks have an aggregate component (i.e., Ξ_t) and an idiosyncratic one (i.e., $\omega_{i,t}$). The aggregate component covaries with financial returns, and, in particular, log income has a loading of β^Y on the robo portfolio's systematic return in year t.

A number of studies find that income skewness improves the performance of life cycle models (e.g., Guvenen, Ozkan and Song (2014); Bagliano, Fugazza and Nicodano (2018); Catherine (2020)). Following Carroll (1997) and Cocco, Gomes and Maenhout (2005), we incorporate skewness by introducing a disaster state in which households receive zero labor income for one year. Such disasters occur with probability ϕ. In years without a disaster, labor income is given by equation (9).

Our empirical analysis primarily concerns investment prior to retirement, and so we model the post-retirement period more simply than do benchmark models. In particular, households do not receive labor income for $t > \underline{T}$. Thus, for $t = \underline{T}+1, ..., \overline{T}$, they solve an "eat-the-pie" problem in which they allocate their liquid assets at retirement, $W_{\underline{T}+1}$, between consumption and savings in the risk-free asset. Appendix C states the problem formally. A household's expected lifetime utility as of age $\underline{T}+1$ then has the familiar form

$$V_{\underline{T}+1}(W_{i,\underline{T}+1}) = B\frac{{W_{i,\underline{T}+1}}^{1-\gamma}}{1-\gamma}, \tag{11}$$

where B is a function of parameters.[12]

6.1.4 Consolidated Problem

In year t, household i allocates shares $\alpha^f_{i,t}$, $\alpha^{\mathcal{S}}_{i,t}$, and $\alpha^{\mathcal{A}}_{i,t}$ of her liquid assets between the risk-free asset, the self-managed portfolio, and the robo portfolio, respectively. She consumes the remaining share $1-\alpha^f_{i,t}-\alpha^{\mathcal{S}}_{i,t}-\alpha^{\mathcal{A}}_{i,t}$,

$$C_{i,t} = [1-\alpha^f_{i,t}-\alpha^{\mathcal{S}}_{i,t}-\alpha^{\mathcal{A}}_{i,t}]W_{i,t}. \tag{12}$$

Thus, the vector $(\alpha^f_{i,t}, \alpha^{\mathcal{S}}_{i,t}, \alpha^{\mathcal{A}}_{i,t})$ defines the problem's control variables. Households optimize over these variables subject to the constraints

$$\alpha^f_{i,t} \geq 0, \quad \alpha^{\mathcal{S}}_{i,t} \geq 0, \quad \alpha^{\mathcal{A}}_{i,t} \geq 0, \tag{13}$$

[12] Explicitly, $B = \sum_{\tau=\underline{T}+1}^{\overline{T}} \delta^{\tau-\underline{T}-1}\left[\frac{1-\chi}{1-\chi^{\overline{T}-\underline{T}}}\right]\left[\delta(1+R^f)\right]^{\frac{\tau-\underline{T}-1}{\gamma}}$ with $\chi = \delta^{\frac{1}{\gamma}}(1+R^f)^{\frac{1-\gamma}{\gamma}}$.

$$1 - \alpha_{i,t}^{f} - \alpha_{i,t}^{\mathcal{S}} - \alpha_{i,t}^{\mathcal{A}} \geq 0, \tag{14}$$

$$\alpha_{i,t}^{\mathcal{A}} = 0 \text{ or } \alpha_{i,t}^{\mathcal{A}} \geq \frac{M}{W_{i,t}}. \tag{15}$$

Constraint (13) rules out borrowing and shorting, which we subsequently relax in Section 8.1. Constraint (14) ensures nonnegative consumption. Both of these constraints are standard. The third constraint, (15), requires an account minimum of M to participate in wealth management.

This setup leads to a problem with two state variables, age (t) and liquid assets ($W_{i,t}$).[13] The latter evolves according to

$$W_{i,t+1} = \left[\alpha_{i,t}^{f}(1 + R^{f}) + \alpha_{i,t}^{\mathcal{S}}(1 + R_{i,t+1}^{\mathcal{S}}) + \alpha_{i,t}^{\mathcal{A}}(1 + R_{i,t+1}^{\mathcal{A}})\right] W_{i,t} + Y_{i,t+1}. \tag{16}$$

Collectively, therefore, household i of age t solves the following Bellman equation,

$$V_t(W_{i,t}) = \max_{\alpha_{i,t}^{f}, \alpha_{i,t}^{\mathcal{S}}, \alpha_{i,t}^{\mathcal{A}}} \left\{ \frac{{C_{i,t}}^{1-\gamma}}{1-\gamma} + \delta p_t \mathbb{E}_t \left[V_{t+1}(W_{i,t+1}) \right] \right\} \tag{17}$$
$$s.t. (8)\text{-}(16).$$

We solve equation (17) using standard numerical methods described in Appendix C. Briefly, we: discretize the state space defined by age and liquid assets; solve equation (17) for age $\underline{T}$; and iteratively solve equation (17) backward for ages $t < \underline{T}$.

6.2 Calibration

Table 5 summarizes the model's parameters and their calibrated values. We first discuss the portfolio parameters and asset pricing factors, shown in panels (a)-(b). Appendix C has details.

6.2.1 Portfolio Parameters

We use the portfolio dataset described in Section 3.2 to realistically calibrate the vector of portfolio parameters, $\left\{\sigma_{\epsilon,i}^{\mathcal{S}}, \sigma_{\epsilon,i}^{\mathcal{A}}, \beta_i^{\mathcal{S}}, \beta_i^{\mathcal{A}}\right\}$. Recall that this dataset includes security-level information on self-managed and counterfactual robo portfolios for households on the margin of participating with the robo advisor, regardless of whether they actually participate.

[13] As noted by Cocco, Gomes and Maenhout (2005), the problem is homogeneous in permanent labor income, $\nu_{i,t}$, allowing us to remove it from the set of state variables.

Our calibration proceeds in three steps. First, we estimate factor loadings and idiosyncratic volatilities for all the individual securities (e.g., stocks, ETFs) in the portfolio dataset. We do so using standard methods in the empirical household finance literature (e.g., Calvet, Campbell and Sodini (2007); Von Gaudecker (2015)), as described in Appendix D. Briefly, for a given vector of risk factors, F, we estimate the following pricing equation for each security k in the portfolio dataset,

$$R_{k,m} = \beta_k F_m + \epsilon_{k,m}, \tag{18}$$

where m indexes month; and $R_{k,m}$ denotes the monthly return on security k in excess of the risk-free return.

We specify F as the three Fama and French (1993) factors with two additional bond factors: the excess returns on U.S. and global bonds, based on Barclays' aggregate bond indices. This 5×1 vector includes many of the factors to which the robo advisor claims to give exposure. Therefore, it likely describes the true return structure more accurately than, say, the CAPM, which we nevertheless consider for robustness. Indeed, Appendix Table A6 shows how robo portfolios have greater net exposure to bonds and to value stocks than their self-managed match. We calibrate the mean and covariance matrix of F using the longest available time series over 1960-2017 and report these values in Appendix Table A7. For reference, the mean and volatility of the market factor equal 7.6% and 14.7%, respectively. Calibrating the factor moments to their long-run values may understate the reduction's welfare impact, given evidence of structurally lower volatility since 2010 (e.g., Smith and Timmermann (2021)).

In the second step, we calculate the parameter vector $\left\{\sigma^{\mathcal{S}}_{\epsilon,j}, \sigma^{\mathcal{A}}_{\epsilon,j}, \beta^{\mathcal{S}}_j, \beta^{\mathcal{A}}_j\right\}$ for each of the 1,913 pairs of self-managed and robo portfolios in the dataset, indexed by j. Let $w^{\mathcal{S}}_j$ denote a vector of weights across securities k for the self-managed portfolio j, and, likewise, let $w^{\mathcal{A}}_j$ denote the weight vector for j's matched robo portfolio. Then, given an estimated vector of loadings across securities, $\hat{\beta}$, and covariance matrix of idiosyncratic volatilities, $\hat{\Sigma_\epsilon}$, we can calculate the portfolio parameters as

$$\sigma^{\mathcal{P}}_{\epsilon,j} = \sqrt{w^{\mathcal{P}\prime}_j \hat{\Sigma_\epsilon} w^{\mathcal{P}}_j}, \quad \beta^{\mathcal{P}}_j = w^{\mathcal{P}\prime}_j \hat{\beta}, \tag{19}$$

for $\mathcal{P} \in \{\mathcal{S}, \mathcal{A}\}$.

Table 6 summarizes the self-managed and matched robo portfolios in our data. Self-managed portfolios are much less diversified than their robo match. In particular, robo portfolios feature a

30 pps higher Sharpe ratio, which, interestingly, holds for both the middle and upper classes. The higher Sharpe ratio partly reflects a 2 pps higher expected return on robo portfolios. By construction, this higher expected return stems from exposure to priced risk. However, robo portfolios contain less total risk because their idiosyncratic volatility is 11 pps lower. The higher idiosyncratic volatility in self-managed portfolios reflects how they concentrate most of their value in individual stocks and actively-managed funds: the median self-managed portfolio allocates only 8% to broad-based index-linked ETFs, versus 100% for robo portfolios. Appendix Table A8 conveys similar patterns based on other factor models.

We also summarize the differences between the self-managed and robo portfolios by performing the return-loss calculation of Calvet, Campbell and Sodini (2007). For each self-managed portfolio, we first calculate the robo-Sharpe ratio expected return which is the product of the robo portfolio's Sharpe ratio and the matched self-managed portfolio's total volatility. Intuitively, this gives a counterfactual self-managed portfolio return under the assumption that the Sharpe ratio is the same as for the robo match. We find that the robo-Sharpe ratio expected return equals nearly 15 pps for both middle and upper class households. Since the actual expected return equals 8 pps, investors lose 7 pps due to much lower mean-variance efficiency when investing on their own.

Lastly, in the third step, we embed the empirical portfolio parameters in the model. Since age and liquid assets comprise the model's state variables, we project each of the empirical portfolio parameters on the holder's log liquid assets and a fifth-order polynomial in age. Then, we substitute the fitted values into the parameter vector $\left\{\sigma_{\epsilon,i}^{\mathcal{S}}, \sigma_{\epsilon,i}^{\mathcal{A}}, \beta_i^{\mathcal{S}}, \beta_i^{\mathcal{A}}\right\}$, using the age and liquid assets of household i. We find similar results when simply substituting the average within wealth-by-age bins in the portfolio dataset.

6.2.2 Other Parameters

We choose preference parameter values of $\gamma = 4$ and $\delta = 0.95$, both of which are fairly standard. Appendix Table A9 obtains similar values when structurally estimating these parameters. We follow Cocco, Gomes and Maenhout (2005) in our calibration of labor income parameters. Accordingly, the deterministic component of income, f_t, is a third-order polynomial in household age, and the coefficients equal those estimated by Cocco, Gomes and Maenhout (2005) for their baseline analysis. Similarly, we parameterize $\sigma_\nu = 0.103$, $\sigma_\xi = 0.271$, and $\beta^Y = 0.001$. Following Carroll (1997), we parameterize $\phi = 0.12\%$ to match the share of households with zero total income, based on the 2016 SCF. The remaining parameter values are: $R^f = 0.2\%$, corresponding to

the average one-month Treasury yield over 2010-2017; $t_0 = 25$, $\underline{T} = 65$, and $\overline{T} = 100$, all of which are standard; and p_t, which we calculate using the Center for Disease Control's mortality tables (Xu et al. (2020)).

7 Rationalizing the Reduced-Form Evidence

Our first set of model results are descriptive. We ask whether a population of rational households solving equation (17) would respond to the reduction in the same way as the households in our quasi-experiment. As a necessary first step, we describe households' optimal investment decisions, or policy functions, since these mechanics drive the model (7.1). Then, we use the model to explain the empirical effect of the reduction (7.2).

7.1 Policy Functions

The model reproduces many common implications of life cycle models, but it also generates several realistic implications that arise due to our unique setting. We succinctly summarize these implications here. Appendix C.2 contains more detail.

As in benchmark models, summarized by Gomes (2020), our model generates the following patterns. Young households households accumulate a buffer stock of liquid assets to insure against idiosyncratic income shocks. Consequently, the savings rate initially decreases in age. Similarly, less-wealthy households exhibit a lower savings rate. As households approach retirement, this pattern reverses and the savings rate begins to increase. Young households also hold a large share of their total wealth in bond-like human capital, and so they invest a higher share of their liquid wealth in risky assets than do middle-age households. Lastly, since we do not incorporate stock market participation costs until Section C.4, almost all households hold risky assets, either through the robo portfolio or through the self-managed portfolio.

Unlike benchmark models, our model features two portfolios, one of which, the robo portfolio, has better diversification, more exposure to priced risk factors, and an account minimum. These features generate a jump in both the robo share and, consequently, the risky share, once the household's liquid assets cross some threshold. Under a $5,000 account minimum, the threshold ranges between $25,000 and $30,000 for most of the household's working life. Notably, this threshold exceeds the actual minimum by a factor of five, which helps explain why we observe large empirical effects for households from the third quintile of the U.S. wealth distribution. This result is

also broadly consistent with our findings in Table 1. For example, the median liquid assets of the existing middle class robo participants are $25,000.

As households approach retirement, the threshold falls for two reasons. First, as in benchmark models, older households consume less and, thus, are willing to allocate a larger share of their consumable resources to the robo advisor. Second, unlike benchmark models, robo portfolios take less priced risk both as the household ages and following declines in wealth, as documented in Figure 6. This glide path strengthens older households' demand for the robo portfolio, even as their declining labor income weakens it.

7.2 Replicating the Effect of the Reduction

7.2.1 Methodology

We replicate the reduction using the policy functions just described. First, we calculate how a household of age t with liquid assets $W_{i,t}$ optimally invests under the previous minimum of $5,000 and, again, under the reduced minimum of $500. Then, using this change in household-level investment, we compare the reduction's theoretical effect on the wealth distribution of robo participants with its empirical effect. We calculate theoretical effects by averaging investment decisions across cells in the grid of age-by-wealth that defines the state space, weighting by the share of households in the 2016 SCF in each cell. Thus, our methodology begins with the acutal U.S. wealth distribution, not a wealth distribution obtained from simulating the model. This distinction matters, since the average household in simulated portfolio choice models rapidly accumulates wealth (e.g., Gomes (2020)). Thus, the reduction would only affect the youngest households, as most older households would have already safely overcome the minimum. In reality, however, households accumulate wealth more slowly, with, for example, at least half of U.S. households over age 40 still within our definition of middle class (Appendix Figure A5). Our methodology accounts for this empirical regularity.

7.2.2 Results

Figure 5 reproduces the empirical democratization of the robo wealth distribution documented in Figure 2. We plot the share of robo participants from each quintile of the U.S. distribution of liquid assets, both for participants who optimally invest under the previous minimum (Existing Participants) and for those who optimally invest under the reduced minimum (New Participants).

The theoretical effect matches its empirical analogue quite well, despite the model's parsimony. By construction, the theoretical effect works through a relaxation of minimum-account constraints, which, per the evidence in Section 5.1, is the channel that drives the empirical effect. We could potentially improve the fit by incorporating attention, peace-of-mind, or trust into the utility function.[14] However, we prefer to err on the side of parsimony, since the spirit of our exercise is to ask how far canonical modelling decisions go toward explaining the empirical effect.

This finding rationalizes the reduction's effect in the following sense: conditional on not investing as well as the robo advisor, middle-class households optimally seek automated wealth management when it becomes more accessible. Weighing in on the origins of household investment inefficiency lies beyond the scope of this paper. It is quite plausible that some form of bounded rationality leads to this inefficiency, but our analysis suggests that households are aware of their own limitations and act rationally in response. This interpretation accords with industry surveys, which report improved asset allocation as a major reason why households seek robo advisors (Costa and Henshaw (2022)).

While our main focus is a change in robo participation, we also examine whether the model can match other features of the data. These tests set a higher bar on our calibration approach. For example, one major concern is that our results on change on robo participation are driven by the choice of model parameters or our specific modeling approach. However, we can gain more confidence in our procedure if the model can also fit other important characteristics of the data.

Table 7 reproduces three other sets of statistics. Panel (a) reports the pre-reduction and post-reduction share of households from the middle-class, which simply aggregates Figure 5 across the two classes. Again, the model fits the pre-reduction share very well. It overshoots the post-reduction share because it fails to match an empirical increase in the share of households from the fourth wealth quintile, which likely reflects factors that lie outside the model. Panel (b) reports the share of liquid assets invested with the robo advisor for new and existing middle-class participants. The model replicates these intensive margin statistics fairly well, although it produces a counterfactually low robo share for new middle-class participants. In the model, these households do not participate before the reduction precisely because they have a low optimal robo share. Lastly, panel (c) uses the extension with stock market participation costs in Section C.4 to calculcate the model-implied stock market participation rate, which matches the empirical rate

[14] For example, including an advertising shock may correct the undershooting of the effect on the fourth quintile and the complementary overshooting of the effect on the second quintile.

very well. The combined results suggest the model quantitatively matches other many features of the data, despite its parsimony

8 Welfare Implications

Our second set of model results concern welfare. We focus on the channels through which the reduction improves welfare and the distribution of this gain across households. As standard, we measure household i's welfare gain by the percent increase in annual consumption under the previous minimum that raises her expected lifetime utility by the same amount as the reduction. Letting $\underline{V}_i$ and $\overline{V}_i$ denote the value functions under the minimums of \$5,000 and \$500, respectively, this welfare gain equals

$$q_i = \left(\frac{\overline{V}_i}{\underline{V}_i}\right)^{\frac{1}{1-\gamma}} - 1, \tag{20}$$

as shown in Appendix C. Table 8 reports various calculations of equation (20), as now described.

8.1 Distributional Effects of the Reduction

Averaging across all robo participants who become participants after the reduction, the welfare gain equals 0.77% of lifetime consumption. This increase constitutes a moderate welfare gain, by the standards of the literature. To further place this increase in perspective, we consider a comparison shock to a household's permanent labor income, while holding the account minimum fixed at \$5,000. Accordingly, we find that a welfare gain of 0.77% approximately maps to a 5 pps increase in permanent income. We consider a similar shock that permanently raises the equity premium, finding that the reduction improves welfare by as much as around a 1.5 pps higher equity premium. These exercises are summarized in panel (d) of Table 9.

There is, however, substantial variation in these welfare gains. Following other papers that use quantitative models to study distributional effects (e.g., Gete and Zecchetto (2018)), we recalculate equation (20) separately for subpopulations defined by wealth and age. Across the wealth distribution, the reduction only benefits middle-class households (column (1)), with no effect on the upper class (column (5)). Across the age distribution, new participants older than 55 gain almost three times as much (1.68%) as new participants aged 35 or younger (0.58%). This surprising finding challenges robo advisors' popular image and their claim to "build our products and services

for millennials" (e.g., Hutchins (2020)).

To understand why the reduction's welfare effect varies by age, first recall that all gains reported in Table 8 condition on households who do not participate under a \$5,000 minimum. Next, consider how many years it would take for such a non-participant to accumulate enough wealth that they decide to participate, again under a \$5,000 minimum. We pursue this question by simulating 5,000 life cycles of robo participation from various starting ages. Then, we calculate the average number of years until the household's first robo investment across simulations. We normalize this average by the remaining years in the household's working life, since our welfare statistic q_i is also defined on an annual basis. The starting level of liquid assets equals \$6,000 in each simulation, the 40^{th}percentile of the U.S. wealth distribution. The starting level of permanent income $\nu_{i,t}$ equals the average value among households who have accumulated less than \$6,000 by the time they reach the starting age.

Figure 6 reports the results. A 55 year-old who does not participate under a \$5,000 minimum would continue to spend around half of her remaining working life as a robo non-participant. By contrast, a similar 25 year-old would spend around 3% of her life as a non-participant. The difference reflects how the 55 year-old has declining labor income, and, moreover, has drawn an unfavorable sequence of permanent income shocks. Put simply, the 55 year-old has strong demand for the robo portfolio, per the discussion in Section 7.1, but she lacks the resources to invest because she has "fallen on tough times". We reiterate how, in the overall U.S. population, households with low-to-moderate levels of wealth actually comprise a large share of households past their peak earnings years (Appendix Figure A5).

Additionally, our simulations suggests that the average change in the probability of robo participation equals 20 pps for households in their 20s, 30s, and 40s, versus 19 pps for households in their 50s and 60s. These results are thus consistent with our earlier reduced-form findings. In particular, column (3) of Table 2 shows that the effect of the reduction in account minimum on middle-class households' robo participation does not significantly vary with age.

8.2 Understanding the Channels: Gain under Counterfactual Robo Portfolios

By construction, any change in welfare must reflect a difference between robo and self-managed portfolio characteristics. We would like to provide some ranking of each characteristic's welfare impact; doing so teaches us about the sources of middle-class households' demand for wealth management. Accordingly, we recalculate the welfare gain from reducing the minimum to \$500

on a counterfactual robo portfolio that no longer differs from the self-managed portfolio in terms of a given characteristic (e.g., uncompensated risk). The difference between the resulting gain and the baseline gain in panel (a) tells us how important that characteristic is, from a welfare perspective.

First, we consider differences in asset allocation, as parameterized by $\sigma^{\mathcal{P}}_{\epsilon,i}$ and $\beta^{\mathcal{P}}_i$. Panel (b) of Table 8 reports the results. Focusing on column (1), we obtain a substantially smaller welfare gain of 0.23% if robo portfolios provide the same level of uncompensated risk as self-managed portfolios ($\sigma^{\mathcal{A}}_{\epsilon,i} = \sigma^{\mathcal{S}}_{\epsilon,i}$). Thus, households value the diversification provided by the robo advisor. Similarly, we obtain a gain of 0.2% if robo portfolios provide the same bond exposure as self-managed portfolios. This finding reflects how robo portfolios provide greater net bond exposure than portfolios managed by households themselves. Since bonds have a relatively high Sharpe ratio (e.g., Asness, Frazzini and Pedersen (2012)), higher bond exposure improves welfare.[15] By contrast, constraining robo portfolios to provide the same stock market exposure (i.e., beta) as self-managed portfolios results in almost the same welfare gain as without this constraint (0.72%). This finding reflects how self-managed and robo portfolios already have similar stock market exposure. Lastly, removing the difference in value (i.e., HML) and size (i.e., SMB) exposure results in a somewhat smaller gain of 0.46%, primarily because self-managed portfolios have less exposure to value stocks. Appendix Table A6 substantiates this paragraph by reporting factor loadings on self-managed and robo portfolios.

Next, we consider how the robo portfolio adjusts according to a household's age and wealth, as already documented in Figure 4. Explicitly, we constrain the robo portfolio's asset allocation: to vary by wealth, but not by age; to vary by age, but not by wealth; and, finally, to not vary at all. Panel (c) reports the welfare gain from reducing the minimum on these counterfactual robo portfolios. We find a smaller gain of 0.37% when removing the adjustment for age, that is, the "glide path". Interestingly, we also find a smaller gain of 0.34% when removing the adjustment for wealth. To interpret, middle-class households value how the robo advisor reduces the portfolio's risk after negative wealth shocks, which most TDFs do not do.[16] Older middle-class households,

[15] This observation undergirds the well-known strategy of risk parity. We clarify that returns follow a stationary distribution in our model. Thus, bond exposure does not add value through reinvestment motives described by Campbell and Viceira (2002).

[16] Personalization by wealth and risk attitude can improve welfare because the two-fund theorem (Tobin (1958)) does not hold in our setting. One reason concerns labor income (e.g., Campbell and Viceira (2002)). To see why, write the effective return on liquid wealth for a robo participant as $\alpha^{\mathcal{A}}_{i,t}(\beta^{\mathcal{A}}_i F_{t+1} + \epsilon^{\mathcal{A}}_{i,t+1}) + \alpha^{f}_{i,t} R^f + \frac{Y_{i,t+1}}{W_{i,t}}$. Since $\frac{Y_{i,t+1}}{W_{i,t}}$ varies by both wealth and age, per equation (9), the optimal choice of $\beta^{\mathcal{A}}_i$ will similarly vary. Intuitively, labor income functions like an endowed risky asset. Since $\beta^{\mathcal{A}}_i$ encodes asset allocation within the robo portfolio, its optimal value will generically

in particular, value this adjustment, as they experience the largest drop in welfare when removing it.

Our counterfacutal analysis suggests that middle-class households value both basic characteristics of robo portfolios and also characteristics that go beyond the basic recommendation to buy a stock index fund. In terms of basic characteristics, robo portfolios add value through reducing uncompensated risk, as would, presumably, an index fund. However, households place the same value on arguably more-sophisticated characteristics, like exposure to additional priced risk factors and personalization by age and wealth. Anecdotally, this finding accords with the advisor's mantra that "everyone deserves sophisticated investment advice".

8.3 Extensions and Robustness

8.3.1 Extensions

Table 9 recalculates the welfare gain under various model extensions, the details of which are in in Appendix C.

The first extension introduces a per-period cost of holding the self-managed portfolio. As already noted, quantitative models generate counterfactually high stock market participation rates without such a cost. In this extension, households incur a per-period cost of $\kappa = \$200$ when holding the self-managed portfolio, which we obtain from the structural estimation described in Appendix C.5. This cost equals around 20% of the inflation-adjusted cost in Vissing-Jørgensen (2002), which reflects how we already incorporate features that discourage stock market participation, like inefficient self-managed portfolios. Under this extension, we calculate a welfare gain of 1.23% due to the reduction, as shown in panel (a). In relation to Table 8, our baseline effect of 0.77% equals around two-thirds of the effect with stock market participation costs. Thus, in a more realistic setting with limited stock market participation, the reduction adds value by bringing households into the stock market, but this channel does not drive the welfare gain. Intuitively, the welfare loss from stock market non-participation is small if the alternative is investing inefficiently on one's own (e.g., Calvet et al. (2007)). Consistent with this interpretation, we find almost the same effect under a higher cost of $\kappa = \$700$.

depend on the distribution of $\frac{Y_{i,t+1}}{W_{i,t}}$. Additionally, the two-fund theorem does not necessarily hold because the account minimum and the implicit borrowing constraint (13) place bounds on the risky share. Lastly, while each household has an optimal robo allocation, the actual robo allocation used to calibrate $R^A_{i,t}$ only approximates the optimal one. Personalization raises welfare insofar as it moves a household closer to her optimal allocation.

Next, we allow households to borrow at the average interest rate on credit card debt in 2015 (12%). The resulting welfare gain in panel (a) equals 0.8%, almost equivalent to the baseline gain. This similarity partly reflects a precautionary savings motive, which discourages middle-class households from borrowing. Moreover, borrowing to invest in the robo portfolio would be sub-optimal given the high interest rate on credit card debt.

In our final extension, we introduce a defined contribution plan. Like in Campbell et al. (2001), households must allocate 10% of their annual income to this plan, which accrues that the risk-free rate, and they cannot withdraw funds until retirement. Thus, a defined contribution plan functions similarly to forced savings in bonds. In this sense, it may increase older households' demand for the robo portfolio. On the other hand, such a plan may reduce older households' relatively high savings rate, thus reducing their demand for the robo portfolio.

8.3.2 Robustness to Miscalibration

We recalculate welfare implications under two different calibrations of portfolio parameters. First, we adjust for selection into consulting an online financial advisor using the Heckman (1979) selection model from Section 5.4. Explicitly, when projecting portfolio characteristics onto the model's state variables as in Section 6.2, we control for the inverse Mills ratio associated with the probability of seeking financial advice online. Second, we calibrate the portfolio parameters using the relationship between the portfolio characteristics of Swedish households and their demographics as estimated in Table 5 of Calvet, Campbell and Sodini (2007). In particular, we use their estimated coefficients to describe a self-managed portfolio's market beta and its idiosyncratic volatility as functions of the model's state variables. Panel (b) shows that the welfare gains under these alternative calibrations are comparable to the baseline.

Lastly, we reproduce the main results under the following parameterizations: structurally estimated preference parameter values; a high discount factor of $\delta = 0.85$; and a 20% correlation between labor income and financial returns. Panel (c) shows that the distribution of welfare gains is robust to these parameterizations. The level of this gain rises under the structurally estimated parameter values, which is driven by a higher coefficient of relative risk aversion ($\gamma = 5.1$).

9 Conclusion

We draw two conclusions. First, from the perspective of the wealth management industry, we

show that wealth management is not necessarily a luxury good. In fact, middle-class households have substantial demand for it. This conclusion comes from studying a large and unexpected reduction in account minimum by a major U.S. automated wealth manager, or robo advisor. The reduction significantly increases the number of robo participants from the middle segments of the U.S. wealth distribution. Interestingly, we document substantial demand from older middle-class households, and we find that households value aspects of wealth management that go beyond simply investing a client's portfolio in a stock index fund (e.g., exposure to additional risk factors, age and wealth glide paths). As the market for automated wealth management grows, either independently or within the banking sector, it may substitute for government programs that aim to provide the less-wealthy with investment opportunities (e.g., myRA, NEST). However, a rigorous treatment of policy implications deserves its own paper.

Second, from the perspective of economic theory, our results support models of bounded rationality in which households act optimally given limits on their ability to invest efficiently on their own. We arrive at this conclusion by quantitatively explaining the previous quasi-experiment with a portfolio choice model calibrated to match portfolio-level data. Households optimally seek professional management because they cannot diversify away uncompensated risk as well as a professional manager can. By reducing its minimum, the automated wealth manager enables modestly wealthy households to benefit from professional management, thus improving their welfare. The overall welfare gain is modest, but there is substantial variation. We leave open the question of why households invest inefficiently on their own, as well as whether these results extend to hybrid wealth management that does not rely fully on automation.

References

Asness, C. S., Frazzini, A. and Pedersen, L. H.: 2012, Leverage Aversion and Risk Parity, *Financial Analysts Journal* . 8.2

Bachas, P., Gertler, P., Higgins, S. and Seira, E.: 2018, Digital Financial Services Go a Long Way: Transaction Costs and Financial Inclusion, *AEA Papers and Proceedings* . 4

Bachas, P., Gertler, P., Higgins, S. and Seira, E.: 2020, How Debit Cards Enable the Poor to Save More, *The Journal of Finance* . 4

Bagliano, F. C., Fugazza, C. and Nicodano, G.: 2018, A Life-Cycle Model with Unemployment Traps. 6.1.3

Balduzzi, P. and Reuter, J.: 2019, Heterogeneity in target date funds: Strategic risk-taking or risk matching?, *The Review of Financial Studies* . 1

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* . 1

Bartlett, R., Morse, A., Stanton, R. and Wallace, N.: 2021, Consumer-Lending Discrimination in the FinTech Era, *Journal of Financial Economics* . 4

Bhattacharya, U., Hackethal, A., Kaesler, S., Loos, B. and Meyer, S.: 2012, Is Unbiased Financial Advice to Retail Investors Sufficient? Answers from a Large Field Study, *Review of Financial Studies* . 2.2

Bianchi, M. and Briére, M.: 2020, Robo-Advising for Small Investors. 1

Bombardini, M. and Trebbi, F.: 2012, Risk Aversion and Expected Utility Theory: An Experiment with Large and Small Stakes, *Journal of the European Economic Association* . 1, B.5

Brennan, M. J., Li, F. and Torous, W. N.: 2005, Dollar Cost Averaging, *Review of Finance* . 3.4

Bricker, J., Dettling, L. J., Henriques, A., Hsu, J. W., Jacobs, L., Moore, K. B., Pack, S., Sabelhaus, J., Thompson, J. and Windle, R. A.: 2017, Changes in U.S. Family Finances from 2013 to 2016: Evidence from the Survey of Consumer Finances, *Federal Reserve Bulletin* . 3.3, A.3.1

Calvet, L. E., Campbell, J. and Sodini, P.: 2007, Down or Out: Assessing the Welfare Costs of Household Investment Mistakes., *Journal of Political Economy* . 1, 2.2, 6.2.1, 6.2.1, 8.3.1, 8.3.2, 9, D, D

Calvet, L. E., Campbell, J. Y. and Sodini, P.: 2009, Fight or Flight: Portfolio Rebalancing by Individual Investors, *The Quarterly Journal of Economics* . 2.2

Campbell, J. Y., Cocco, J. a. F., Gomes, F. J. and Maenhout, P. J.: 2001, Investing Retirement Wealth: A Life-Cycle Model, *Risk Aspects of Investment-Based Social Security Reform*. 6.1, 8.3.1, 9, C.4.3

Campbell, J. Y. and Viceira, L. M.: 2002, *Strategic Asset Allocation: Portfolio Choice for Long-term Investors*, Oxford University Press. 15, 16

Carroll, C. D.: 1997, Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis, *The Quarterly Journal of Economics* . 6.1.3, 6.1.3, 6.2.2, 5

Catherine, S.: 2020, Countercyclical Labor Income Risk and Portfolio Choices over the Life-Cycle. 6.1.3

Chalmers, J. and Reuter, J.: 2020, Is Conflicted Investment Advice Better than No Advice?, *Journal of Financial Economics* . 2, 3.2

Christelis, D., Jappelli, T. and Padula, M.: 2010, Cognitive Abilities and Portfolio Choice, *European Economic Review* . 1, 2.2

Cocco, J. F., Gomes, F. J. and Maenhout, P. J.: 2005, Consumption and portfolio choice over the life cycle, *The Review of Financial Studies* . 1, 6.1, 6.1.3, 13, 6.2.2, 5, C.1.2

Cole, S., Paulson, A. and Shastry, G. K.: 2014, Smart Money? The Effect of Education on Financial Outcomes, *Review of Financial Studies* . 1

Costa-Dias, M. and O'Dea, C.: 2019, Dynamic Economics in Practice, *Institute for Fiscal Studies Working Paper* . C.1.1

Costa, P. and Henshaw, J. E.: 2022, Quantifying the Investor's View on the Value of Human and Robo-Advice, *Vanguard Research Report* . 7.2.2

D'Acunto, F., Prabhala, N. and Rossi, A. G.: 2019, The Promises and Pitfalls of Robo-Advising, *Review of Financial Studies* . 1

D'Acunto, F. and Rossi, A. G.: 2020, Robo-Advising, *in* R. Rau, R. Wardrop and L. Zingales (eds), *Palgrave Macmillan Handbook of Technological Finance*. 1, 2, 2.1

Deaton, A.: 1991, Savings and Liquidity Constraints, *Econometrica* . 6.1.1

D'Hondt, C., De Winne, R., Ghysels, E. and Raymond, S.: 2020, Artificial Intelligence Alter Egos: Who Might Benefit from Robo-Investing?, *Journal of Empirical Finance* . 1

Egan, M.: 2019, Brokers versus Retail Investors: Conflicting Interests and Dominated Products, *The Journal of Finance* . 2

Fagereng, A., Gottlieb, C. and Guiso, L.: 2017, Asset Market Participation and Portfolio Choice over the Life Cycle, *The Journal of Finance* . 1

Fama, E. F. and French, K. R.: 1993, Common Risk Factors in the Returns on Stocks and Bonds, *Journal of Financial Economics* . 6, 6.2.1, 9

Foerster, S., Linnainmaa, J. T., Melzer, B. T. and Previtero, A.: 2017, Retail Financial Advice: Does One Size Fit All?, *The Journal of Finance* . 2

French, K. R.: 2008, Presidential Address: The Cost of Active Investing, *The Journal of Finance* . 1

Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T. and Walther, A.: 2021, Predictably Unequal? The Effects of Machine Learning on Credit Markets, *The Journal of Finance* . 4

Fuster, A., Plosser, M., Schnabl, P. and Vickery, J.: 2019, The Role of Technology in Mortgage Lending, *Review of Financial Studies* . 4

Gârleanu, N. and Pedersen, L. H.: 2018, Efficiently Inefficient Markets for Assets and Asset Management, *The Journal of Finance* . 1

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

Gete, P. and Zecchetto, F.: 2018, Distributional Implications of Government Guarantees in Mortgage Markets, *Review of Financial Studies* . 8.1

Gomes, F.: 2020, Portfolio Choice Over the Life Cycle: A Survey, *Annual Review of Financial Economics* . 1, 7.1, 7.2.1, C.2

Grinblatt, M., Keloharju, M. and Linnainmaa, J.: 2011, IQ and Stock Market Participation, *The Journal of Finance* . 1, 2.2

Guvenen, F., Ozkan, S. and Song, J.: 2014, The Nature of Countercyclical Income Risk, *The Journal of Political Economy* . 6.1.3

Haliassos, M. and Bertaut, C. C.: 1995, Why do so Few Hold Stocks?, *The Economic Journal* . 1

Heckman, J. J.: 1979, Sample Selection Bias as a Specification Error, *Econometrica* . 1, 5.4, 8.3.2, 9, B.3, B.3, A10

Higgins, S.: 2020, Financial Technology Adoption. 4

Hong, C. Y., Lu, X. and Pan, J.: 2020, FinTech Adoption and Household Risk-Taking, *NBER Working Paper* . 4

Hong, H., Kubik, J. D. and Stein, J. C.: 2004, Social interaction and stock market participation, *The Journal of Finance* . 1

Hutchins, C.: 2020, How Millennials Are Thinking About Their Finances During COVID-19, *Technical report*, Wealthfront. 1, 8.1

Jurek, J. W. and Viceira, L. M.: 2011, Optimal Value and Growth Tilts in Long-Horizon Portfolios, *Review of Finance* . 2.2

Kaniel, R. and Parham, R.: 2017, WSJ Category Kings - The Impact of Media Attention on Consumer and Mutual Fund Investment Decisions, *Journal of Financial Economics* . 1

Linnainmaa, J. T., Melzer, B. T. and Previtero, A.: 2021, The Misguided Beliefs of Financial Advisors, *The Journal of Finance* . 2

Loos, B., Previtero, A., Scheurle, S. and Hackethal, A.: 2020, Robo-advisers and Investor Behavior. 1

Malkiel, B.: 2015, *A Random Walk Down Wall Street: the Time-Tested Strategy for Successful Investing*, W.W. Norton and Company, Inc. 2.2

Merton, R. C.: 1980, On Estimating the Expected Return on the Market: An Exploratory Investigation, *Journal of Financial Economics* . 6.1.2

Moulliet, D., Stolzenbach, J., Bein, A. and Wagner, I.: 2016, Cost-income ratios and robo-advisory, *Deloitte Technical Report* . 2.1, 2.3

Mullainathan, S., Noeth, M. and Schoar, A.: 2012, The Market for Financial Advice: An Audit Study, *NBER Working Paper* . 2

Newey, W.: 2007, Course Materials for Nonlinear Econometric Analysis, *MIT OpenCourseWare* . C.5

Parker, J. A., Schoar, A., Cole, A. and Simester, D.: 2022, Household Portfolios and Retirement Saving over the Life Cycle, *NBER Working Paper* . 1

Parker, J. A., Schoar, A. and Su, Y.: 2021, Retail Financial Innovation and Stock Market Dynamics: The Case of Target Date Funds, *NBER Working Paper* . 1

Philippon, T.: 2016, The FinTech Opportunity, *NBER Working Paper* . 1

Philippon, T.: 2019, On FinTech and Financial Inclusion, *NBER Working Paper* . 4

Rossi, A. and Utkus, S.: 2020, Who Benefits from Robo-Advising: Evidence from Machine Learning. 1

Securities and Exchange Commission (SEC): 2017, Form ADV and IARD Frequently Asked Questions: Item 5.D, *SEC Reference Guides* .

Smith, S. C. and Timmermann, A.: 2021, Break Risk, *Review of Financial Studies* . 6.2.1

Tauchen, G. and Hussey, R.: 1991, Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models, *Econometrica* . C.1.2

The Economist: 2019/12/18, For the Money, not the Few: Wealth Managers are Promising Business-Class Service for the Masses, *The Economist* .

Thomson Reuters: 2015/07/07, Robo-Advisor Wealthfront Lowers Account Minimum to $500, *Thomson Reuters* .

Tobin, J.: 1958, Liquidity Preference as Behavior Towards Risk, *Review of Economic Studies* **25**. 16

Van Rooij, M., Lusardi, A. and Alessie, R.: 2011, Financial Literacy and Stock Market Participation, *Journal of Financial Economics* . 1, 1

Vissing-Jørgensen, A.: 2002, Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structure, *NBER Working Paper* . 1, 8.3.1

Vissing-Jørgensen, A.: 2003, Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions, *NBER Macroeconomics Annual* . C.4.1

Von Gaudecker, H.-M.: 2015, How Does Household Portfolio Diversification Vary with Financial Literacy and Financial Advice?, *The Journal of Finance* . 1, 6.2.1

Xu, J., Murphy, S. L., Kochanek, K. D. and Arias, E.: 2020, Mortality in the United States, 2018, *NCHS Data Brief* . 6.2.2, A.3.3

Figure 1: Shift in Wealth Distribution of Robo Participants

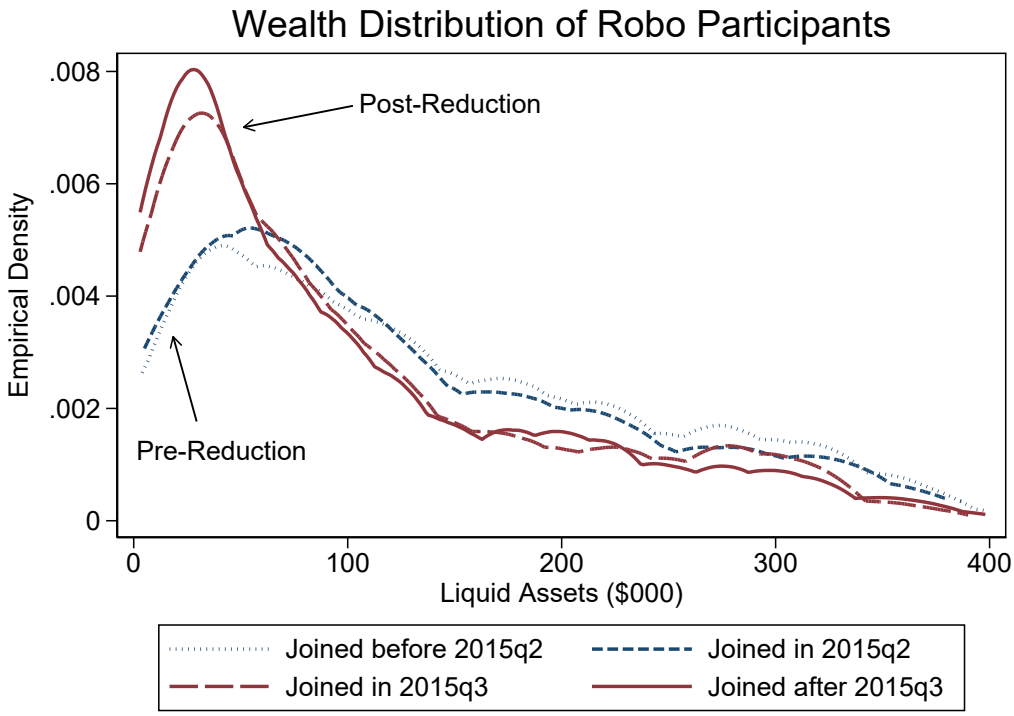

Note: This figure plots the distribution of liquid assets among households who participated with the robo advisor before the reduction in account minimum (Existing Participants) and who become robo participants after the reduction (New Participants). Liquid assets are defined in Table 1. We remove a very small number of households with liquid assets greater than $400,000 so that the distribution can fit on a page. The distribution is calculated using a kernel density.

Figure 2: Change in Representativeness of Robo Wealth Distribution

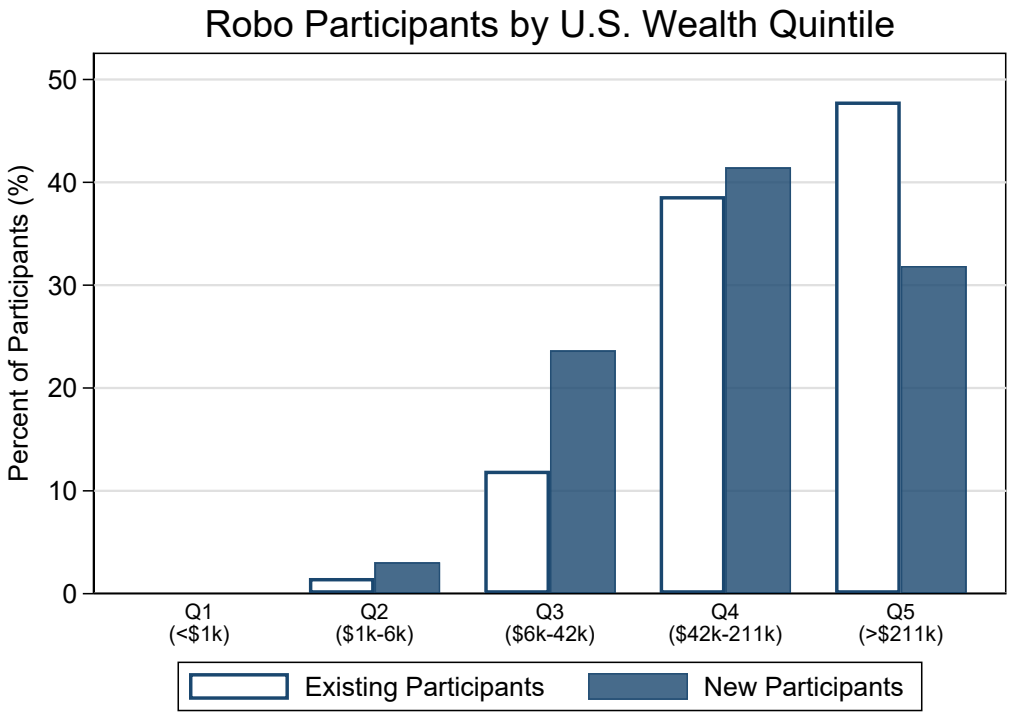

Note: This figure plots the share of robo participants from each quintile of the U.S. wealth distribution. The share is calculated separately for households who participated before the reduction in account minimum (Existing Participants) and who become participants after the reduction (New Participants). Wealth consists of liquid assets, defined in Table 1. Wealth quintiles are calculated using the 2016 Survey of Consumer Finances (SCF) dataset.

Figure 3: Event Study. Pre-Trends in Robo Participation by Wealth Class.

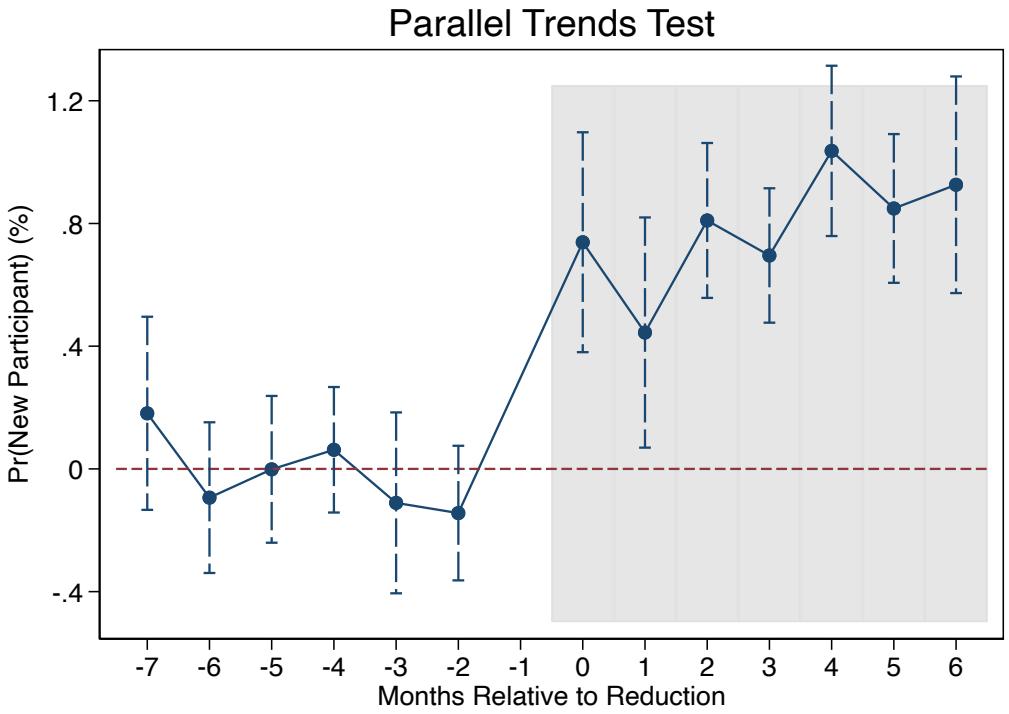

Note: This figure assesses the scope for pre-trends in middle-class households' robo participation by estimating the following specification:

$$New\ Participant_{i,t} = \sum_{m \neq June\,2015} (\mu_m \times Middle_i \times \mathbb{1}_{t \in m}) + \zeta_i + \varrho_t + \sum_{m \neq June\,2015} (\psi_m \times X_i \times \mathbb{1}_{t \in m}) + u_{i,t}.$$

The figure shows the estimated coefficients μ_m,which are interpreted as the average difference in the weekly probability of robo participation between the middle and the upper classes among all weeks in a given month. Brackets are 95% confidence intervals with standard errors clustered by household and week. The shaded region corresponds to the period after the reduction in account minimum. Time 0 is the month when the reduction was implemented. The remaining notes are the same as in Table 2.

Figure 4: Risk Profile of Robo Portfolios

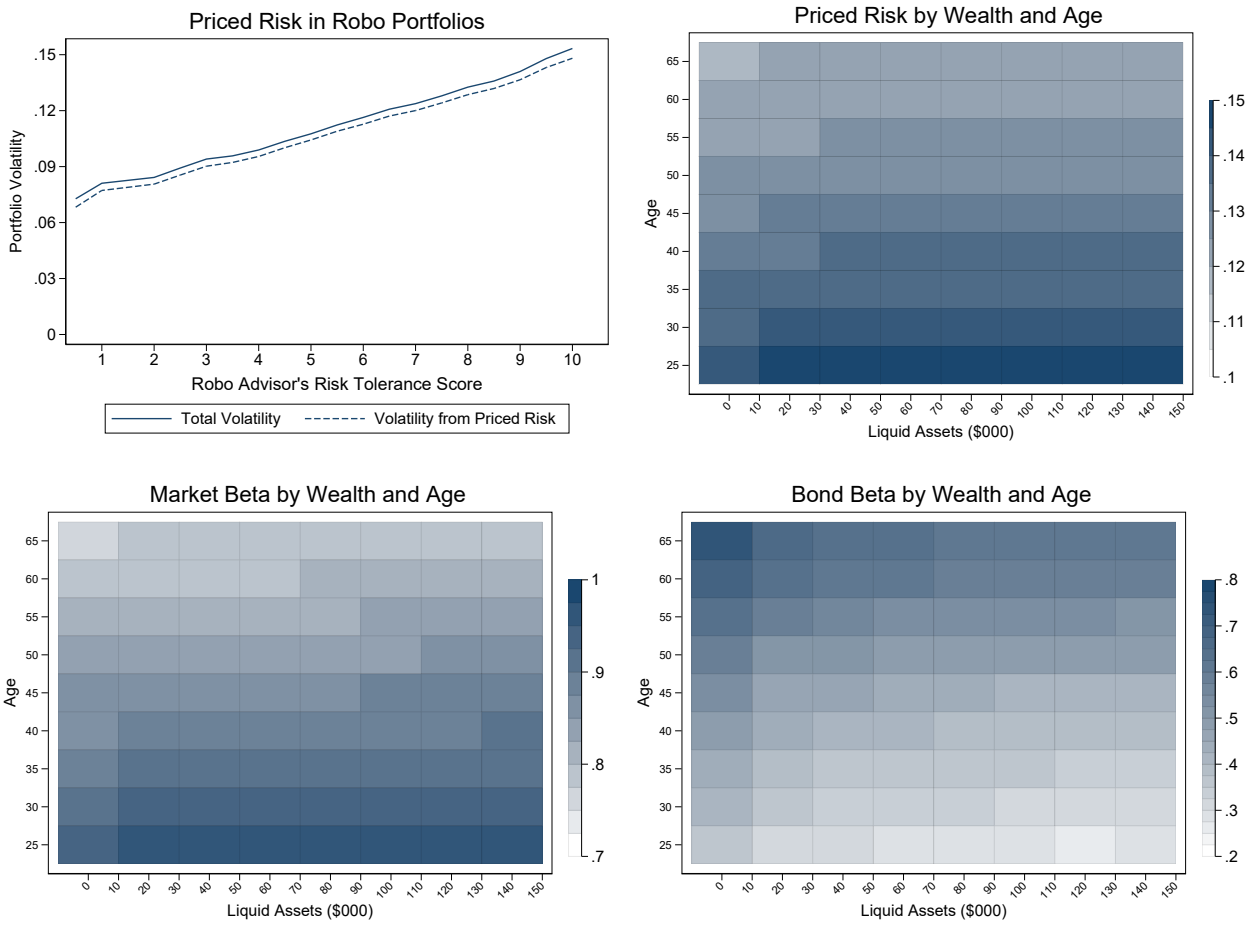

Note: This figure summarizes the priced risk (i.e., systematic risk) in robo portfolios. Panel (a) plots the volatility of total return and priced risk of robo portfolios by the portfolio's risk tolerance score, which the robo advisor assigns based on the household's response to the advisor's questionnaire. *Total Volatility* is the standard deviation of the total portfolio returns. *Volatility from Priced Risk* is the standard deviation of the portion of the portfolio return which is attributed to priced risk factors. Priced risk is measured using a five-factor asset pricing model. The model is estimated at the security level, and the factor loadings are aggregated across securities using the portfolio weights that the robo advisor assigns to households of a given age and liquid wealth. The five factors are: the three Fama and French (1993) factors, and the excess returns on U.S. and global bonds. Priced risk is defined as the volatility of the product of the factor loadings and the factors. Panel (b) plots priced risk by household age and wealth. Panels (c)-(d) plot the loadings on specific factors: the market (*Market Beta*) and U.S. bonds (*Bond Beta*). Details on robo portfolios are in Section 3. Details on estimating the factor models are in Appendix D.

Figure 5: Theoretical Change in Representativeness of Robo Wealth Distribution.

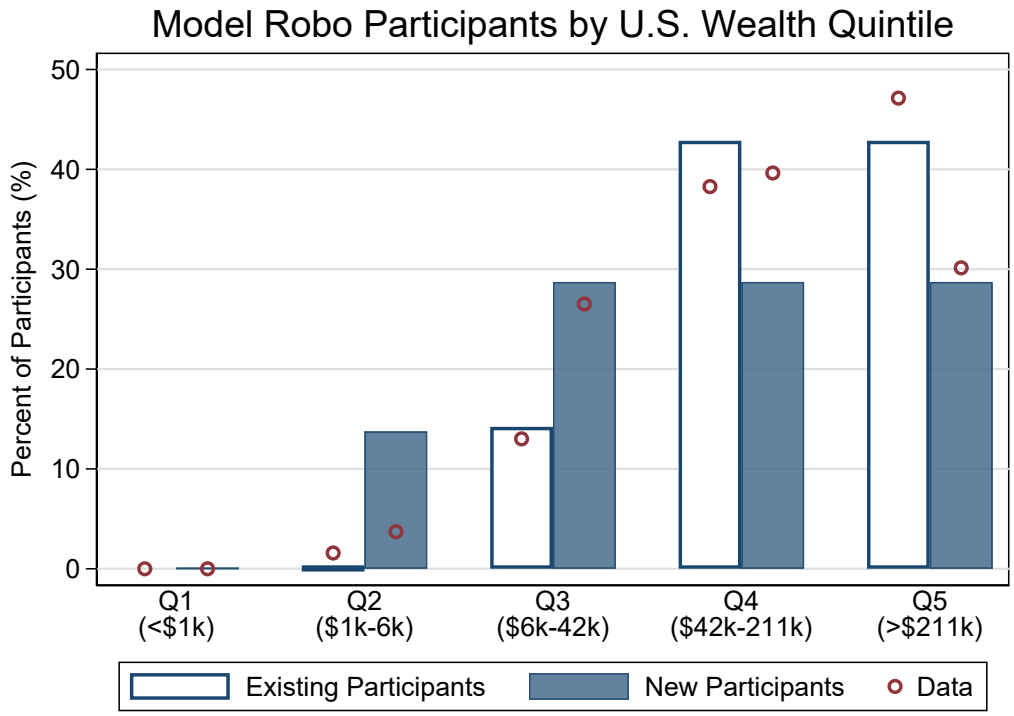

Note: This figure plots the share of robo participants from each quintile of the U.S. wealth distribution, based on the portfolio choice model in Section 6. Explicitly, figure shows the distribution across wealth quintiles for participants who find it optimal to participate under the previous minimum (Existing Participants) and for those who find it optimal to participate under the reduced minimum (New Participants). The overall number of existing and new participants is calculated by aggregating household-level policy functions across the bins of age and liquid assets that define the state space, weighting by the share of households in the 2016 SCF within each bin. The red open circles show the empirical share of robo participants from each quintile of the U.S. wealth distribution based on Figure 2. The remaining notes are the same as in Figure 2.

Figure 6: Simulated Time to First Robo Investment by Age.

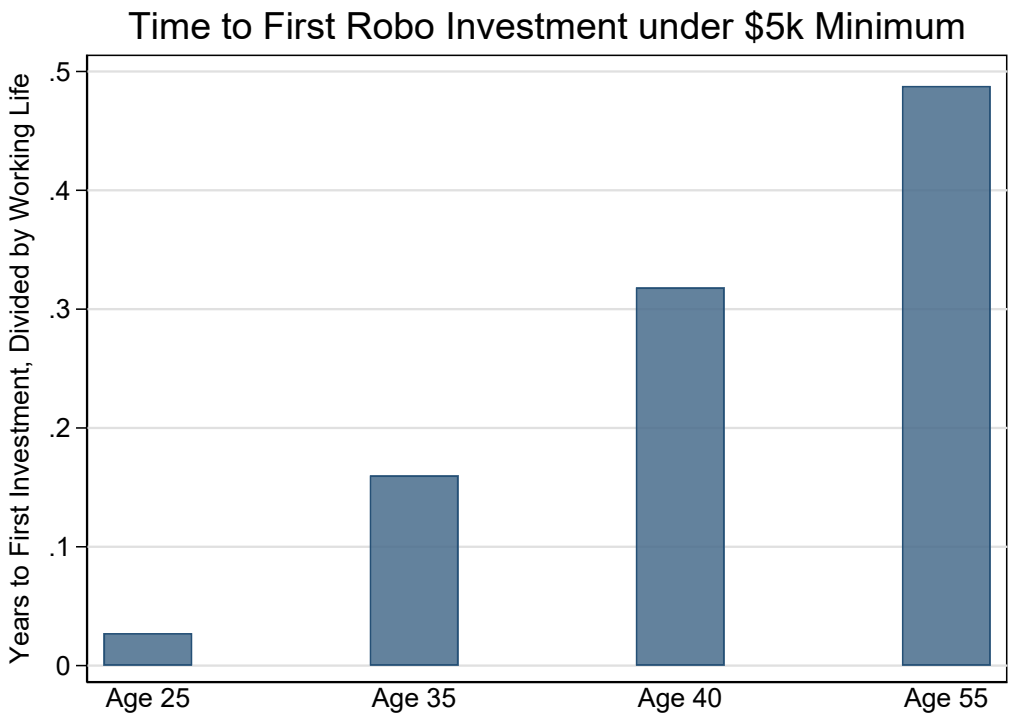

Note: This figure plots the amount of time until a middle-class household first participates with the robo advisor under an account minimum of $5,000. The amount of time is calculated based on 5,000 simulated paths of the model in Section 6. The outcome variable equals the average number of years until a household who begins age t with $6,000 in liquid assets becomes a robo participant, where t equals one of the values shown in the figure. The number of years until participation is then normalized by the number of years until the household retires. The starting level of liquid assets equals the 40^{th}percentile of the U.S. wealth distribution ($6,000), and the starting age equals one of the ages shown in the figure. The starting level of permanent income, $\nu_{i,t}$, is chosen to match the average value of permanent income among households who have reached age t with less than $6,000 in liquid assets. We obtain $\nu_{i,t}$ by separately simulating the life cycle of 5,000 households starting at age 25 and calculating the distribution of $\nu_{i,t}$ at each subsequent age.

Table 1: Summary of Robo Participants.

	Existing Participants ($N = 4,366$)			New Participants ($N = 5,336$)			New vs. Existing	
	Mean	Standard Deviation	Median	Mean	Standard Deviation	Median	Difference in Mean	Standard Error
(a) All Households:								
Liquid Assets$_i$ ('000)	436.44	660.82	200	265.21	480.25	100	-171.22***	(11.968)
Income$_i$ ('000)	157.36	110.67	130	116.17	95.9	90	-41.18***	(2.128)
Initial Deposit$_i$ ('000)	33.68	94.54	10	22.56	72.61	5	-11.12***	(2.424)
Age$_i$	35.79	8.72	34	35.4	9.97	33	-0.39***	(0.190)
Middle$_i$	0.15	0.35		0.3	0.46		0.156***	(0.008)
No Account Closure$_i$	0.95	0.23		0.98	0.15		0.031***	(0.004)
Subsequent Deposit$_i$	0.9	0.3		0.71	0.45		-0.185***	(0.008)
First Deposit at \$5k$_i$	0.33	0.47		0.11	0.32		-0.212***	(0.012)
First Deposit < \$5k$_i$	0	0		0.42	0.49		0.418***	(0.007)
(b) Middle Class:								
Liquid Assets$_i$ ('000)	23.23	11.68	25	19.71	11.36	18	-3.527***	(0.542)
Income$_i$ ('000)	92.86	62.21	80	67.14	42.52	60	-25.720***	(2.682)
Initial Deposit$_i$ ('000)	7.6	5.34	5	4.95	12.58	2	-2.652***	(0.430)
Age$_i$	30.33	6.33	29	30.04	7.07	28	-0.293	(0.306)
No Account Closure$_i$	0.92	0.27		0.97	0.18		0.043***	(0.011)
Subsequent Deposit$_i$	0.86	0.34		0.72	0.45		-0.149***	(0.018)
First Deposit at \$5k$_i$	0.55	0.5		0.09	0.29		-0.460***	(0.028)
First Deposit < \$5k$_i$	0	0		0.66	0.47		0.661***	(0.012)

Note: This table summarizes households who participated with the robo advisor before the reduction in account minimum (Existing Participants) and who become participants after the reduction (New Participants), based on the Deposits Dataset. Subscript i indexes household. *Liquid Assets*$_i$ is the sum of cash, savings accounts, certificates of deposit, mutual funds, IRAs, 401ks, and public stocks, in thousands of dollars. *Income*$_i$ is annual household income, in thousands of dollars. *Initial Deposit*$_i$ is the value of the household's initial deposit, in thousands of dollars. *Age*$_i$ is the householder's age. *High Risk Tolerance*$_i$ indicates if the household chooses a higher risk tolerance score than that recommended by the robo advisor. *Middle*$_i$ indicates if i belongs to the second (\$1k-\$6k) or third U.S. wealth quintile (\$6k-\$42k). *No Account Closure*$_i$ indicates if the household does not close the account over the sample period. *Subsequent Deposit*$_i$ similarly indicates if the household makes a subsequent deposit after her initial one. *First Deposit at \$5k*$_i$ indicates if the household's initial deposit equals \$5,000 or is no more than 5% higher. *First Deposit < \$5k*$_i$ indicates if the household's initial deposit lies between \$500 and \$5,000. Wealth consists of liquid assets, and wealth quintiles are calculated using the 2016 Survey of Consumer Finances (SCF) dataset. The sample consists of households who participate with the robo advisor and make a deposit over the period from December 2014 through February 2016. The upper panel summarizes all households in the sample, and the lower panel summarizes households from the second or third U.S. wealth quintile. Appendix A has details.

Table 2: Effects of the Reduction on Middle-Class Participation.

$Y_i =$	First-Difference						Dynamic DiD	
	New Participant$_i$						*New Participant*$_{i,t}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Middle$_i$	15.051***	13.718***	17.987***	8.60**	14.516***	15.467***		
	(1.357)	(1.339)	(5.116)	(3.583)	(3.060)	(1.408)		
Middle$_i \times$ *Age*$_i$			-0.137					
			(0.158)					
Middle$_i \times$ *Post*$_t$							0.823***	0.608***
							(0.067)	(0.085)
Measure of Middle	2nd or 3rd Quintile	2nd or 3rd Quintile	2nd or 3rd Quintile	2nd or 3rd Quintile	2nd Quintile	Buffer on Middle	2nd or 3rd Quintile	2nd or 3rd Quintile
Subsample	Main	Main	Main	Outside Portfolios	Main	Main	Main	Main
Controls	Yes	Yes	Yes	Yes	Yes	Yes		
State FE	No	Yes	Yes	No	Yes	Yes		
Household FE							Yes	Yes
Week FE							Yes	Yes
Controls $\times$ *Post*$_t$							No	Yes
State FE $\times$ *Post*$_t$							No	Yes
R-squared	0.067	0.097	0.097	0.035	0.078	0.098	0.008	0.008
Number of Observations	9,349	9,349	9,349	2,588	7,530	8,982	620,928	620,928

Note: Standard errors are in parentheses. Point estimates are multiplied by 100. This table estimates equations (2) and (4), which test whether the reduction in account minimum brings middle-class households into the market for automated wealth management. Subscripts i and t index household and week. Columns (1)-(6) estimate the first-difference equation (2), and so there are no t subscripts,

$$New\ Participant_i = \mu Middle_i + \psi X_i + \varrho + u_i,$$

where *Middle*$_i$ indicates if i belongs to the second (\$1k-\$6k) or third U.S. wealth quintile (\$6k-\$42k), as opposed to the fourth or fifth quintile (>\$42k) that together constitute the reference group; and *New Participant*$_i$ indicates if i becomes a robo participant after the reduction, as opposed to before it. Columns (5)-(6) assess the scope for measurement error by remeasuring *Middle*$_i$ using alternative measures: an indicator for whether i belongs to the second U.S. wealth quintile, after assigning a missing to households from the third U.S. wealth quintile (Second Wealth Quintile); an indicator for whether i belongs to the second or third U.S. wealth quintile, after assigning a missing value to households whose liquid assets are within a 10% buffer of the third quintile (Buffer on Middle). The sample in columns (1)-(3) and (5)-(6) consists all robo participants in the Deposits Dataset. The sample in column (4) consists of all households in the Portfolio Dataset, which includes robo non-participants and a subset of households in the Deposits Dataset. Wealth consists of liquid assets, defined in Table 1. Wealth quintiles are calculated using the 2016 Survey of Consumer Finances (SCF) dataset. Household controls X_i are: the log of annual household income; the householder's age; and an indicator for whether the household chooses a lower risk tolerance score than that recommended by the robo advisor. State fixed effects cannot be included in column (4) because the state of residence is unobserved for robo non-participants.
Columns (7)-(8) estimate the dynamic difference-in-difference (DiD) baseline equation (4),

$$New\ Participant_{i,t} = \mu\left(Middle_i \times Post_t\right) + \zeta_i + \varrho_t + \psi_t X_i + u_{i,t},$$

where *Post*$_t$ indicates if t is greater than the week of the reduction; and *New Participant*$_{i,t}$ indexes if i becomes a robo participant in week t, as opposed to the other weeks in our observation window. Column (7) includes week and household fixed effects. Column (8) additionally interacts *Post*$_t$ with household controls and a vector of state fixed effects. The point estimates are multiplied by 100 for easier interpretation. Standard errors are clustered by household (heteroskedasticity robust) in columns (1)-(6). Standard errors are two-way clustered by household and week in columns (7)-(8).

Table 3: Testing the Constraints Channel. The Effect on Bunching at the Previous Minimum.

$Y_{i,t} =$	$\mathbb{1}[First\ Deposit_{i,t} = \$5k]$				
	(1)	(2)	(3)	(5)	(6)
$Middle_i \times Post_t$	-30.236***	-30.538***	-23.724***		
	(3.175)	(3.189)	(3.475)		
$Middle_i$	27.129***	24.673***	19.656***		
	(3.028)	(3.131)	(3.273)		
$Lower\ Middle_i \times Post_t$				-43.073***	-36.419***
				(8.400)	(7.837)
$Lower\ Middle_i$				33.922***	29.390***
				(8.135)	(7.586)
$Upper\ Middle_i \times Post_t$				-29.276***	-22.468***
				(3.355)	(3.714)
$Upper\ Middle_i$				23.814***	18.759***
				(3.294)	(3.506)
Week FE	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	Yes	Yes	Yes
State FE	No	Yes	Yes	Yes	Yes
Controls $\times Post_t$	No	No	Yes	No	Yes
State FE $\times Post_t$	No	No	Yes	No	Yes
R-squared	0.113	0.119	0.120	0.281	0.288
Number of Observations	7,162	7,162	7,162	7,162	7,162

Note: Standard errors are in parentheses. Point estimates are multiplied by 100. This table estimates a variant of equation (4) that assesses the robustness of interpreting households from the second or third U.S. wealth quintiles as constrained by the previous account minimum. Subscripts i and t index household and week. The baseline regression equation is of the form,

$$\mathbb{1}[First\ Deposit_{i,t} = \$5k] = \mu\ (Middle_i \times Post_t) + \varrho_t + \psi_t X_i + u_{i,t},$$

where $\mathbb{1}[First\ Deposit_{i,t} = \$5k]$ indicates the initial deposit made by household i in week t equals the previous account minimum or is no more than 5% higher. This variable measures bunching of initial deposits around the previous minimum. The sample consists of household-weeks in which the household makes its initial deposit, and, thus, we cannot include household fixed effects. Columns (4)-(5) estimate separate effects for households from the second U.S. wealth quintile ($Lower\ Middle_i$) and the third quintile ($Upper\ Middle_i$). The remaining notes are the same as in Table 2.

Table 4: Changes in Visibility as a Confounding Factor.

$Y_{i,t} =$	$New\ Participant_{i,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$Middle_i \times Post_t$	0.784***	0.513**	0.487*	0.806***	0.842**	0.817**
	(0.166)	(0.188)	(0.188)	(0.231)	(0.291)	(0.294)
$Middle_i \times Post_t \times Adv_t$	-0.004	0.004	0.004	-0.033	-0.106	-0.106
	(0.013)	(0.013)	(0.013)	(0.083)	(0.102)	(0.102)
$Middle_i \times Adv_t$	-0.006	-0.008	-0.008	-0.026	-0.024	-0.023
	(0.006)	(0.006)	(0.006)	(0.018)	(0.022)	(0.022)
Measure of Advertising (Adv)	Google News			Advisor Blog Posts		
Household FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls $\times Post_t$	No	Yes	Yes	No	Yes	Yes
Controls $\times Adv_t$	No	Yes	Yes	No	Yes	Yes
Controls $\times Adv_t \times Post_t$	No	Yes	Yes	No	Yes	Yes
State FE $\times Post_t$	No	No	Yes	No	No	Yes
State FE $\times Adv_t$	No	No	Yes	No	No	Yes
State FE $\times Adv_t \times Post_t$	No	No	Yes	No	No	Yes
R-squared	0.011	0.012	0.012	0.011	0.012	0.012
Number of Observations	504,504	504,504	504,504	504,504	504,504	504,504

Note: Standard errors are in parentheses. Point estimates are multiplied by 100. This table estimates a variant of equation (4) that assesses the robustness of the baseline results to a dynamic specification that accounts for changes in advertising that may disproportionately affect middle-class households. Subscripts i and t index household and week. The baseline regression equation in columns (1)-(6) is of the form

$$New\ Participant_{i,t} = \mu\left(Middle_i \times Post_t\right) + \lambda_0\left(Middle_i \times Post \times Adv_t\right) + \lambda_1\left(Middle_i \times Adv_t\right) + \zeta_i + \varrho_t + \psi_t X_i + u_{i,t},$$

where Adv_t measures the intensity of the robo advisor's advertising or, more generally, the robo advisor's visibility, in week t. The measures of Adv_t are: the number of news articles about the robo advisor published in the month of week t (Google News); and the number of blog posts written by the robo advisor in the month of week t (Advisor Blog Posts). The news articles used to construct the measure in columns (1)-(3) are the top 150 articles, sorted by relevance, from a Google News search of the advisor's name ("Wealthfront") among articles published in 2015. The remaining notes are the same as in Table 2.

Table 5: Model Parameters

Parameter	Value	Source
(a) Portfolio Parameters:		
Idiosyncratic Volatility ($\sigma_{\epsilon,i}$)	Table 6	Portfolio Dataset
Factor Loadings (β_i)	Appendix Table A6	Portfolio Dataset
(b) Asset Pricing Factors:		
Market Factor, Mean	0.076	CRSP
Market Factor, Volatility	0.147	CRSP
Fama-French Factors	Appendix Table A7	French
Bond Factors	Appendix Table A7	Bloomberg-Barclays
(c) Preferences:		
Coefficient of Relative Risk Aversion (γ)	4	Standard
Discount Factor (δ)	0.95	Standard
(d) Labor Income Parameters:		
Age Profile ($f(x)$)	$0.1682 \cdot x - 0.0323 \cdot x^2/10 + 0.002 \cdot x^3/100$	CGM
Permanent Shock Volatility (σ_v)	0.103	CGM
Temporary Shock Volatility ($\sigma_{\tilde{\xi}}$)	0.271	CGM
Loading on Financial Return (β^Y)	0.001	CGM
Probability of Disaster (ϕ)	0.001	SCF
(e) Other Parameters:		
Risk-Free Rate (R^f)	0.002	French
Pre-Retirement Age Range ($[t_0, \underline{T}]$)	[25, 65]	CDC
Range of Survival Rates ($[p(\underline{T}), p(t_0)]$)	[0.865, 0.999]	CDC

Note: This table summarizes the baseline calibration of the portfolio choice model in Section 6. Panel (a) notes the location of the table summarizing portfolio parameters. Panel (b) summarizes asset pricing factors, presenting the mean and volatility of the market factor and notes the location of the tables summarizing the other factors. Factor moments are calibrated using the means and covariances evaluated over the longest available time series over 1960-2017. Panel (c) summarizes preference parameters. Panel (d) summarizes parameters of the labor income process. Note that: a loading of log labor income on financial returns of $\beta^Y = 0.001$ corresponds to a correlation coefficient of 1%; the probability of labor income disaster is calculated as the share of households with zero total income (Carroll (1997)), using the 2016 SCF. Panel (e) summarizes other parameters: the risk-free rate corresponds to the average one-month Treasury yield in 2016; households begin their problem at age t_0, retire after age $\underline{T}$, and leave the model at $\overline{T} = 100$; and the survival rate corresponds to the probability that a household of age t survives until age $t+1$, and it is monotonically decreasing in age. Column (3) reports the source of each value: CRSP denotes the annually-updated stock file from CRSP; French denotes Ken French's website; Bloomberg-Barclays denotes the Bloomberg-Barclays aggregate U.S. and unhedged global bond indices; CGM denotes Cocco, Gomes and Maenhout (2005); CDC denotes the Center for Disease Control's mortality tables; SCF denotes the 2016 Survey of Consumer Finances; and Portfolio Dataset denotes the paper's portfolio dataset summarized in Table 6. Appendix A has details on these data sources.

Table 6: Summary of Self-Managed and Robo Portfolios

	Middle Class ($N = 354$)			Upper Class ($N = 1,559$)		
	Self-Managed	Matched Robo	Difference	Self-Managed	Matched Robo	Difference
	(1)	(2)	(3)	(4)	(5)	(6)
Sharpe Ratio	0.452	0.750	0.298	0.459	0.756	0.297
			(0.000)			(0.000)
Expected Return	0.080	0.102	0.023	0.078	0.101	0.023
			(0.000)			(0.000)
Total Volatility	0.209	0.137	-0.071	0.196	0.134	-0.062
			(0.000)			(0.000)
Idiosyncratic Volatility	0.146	0.034	-0.111	0.138	0.033	-0.104
			(0.000)			(0.000)
Robo-Sharpe Expected Return	0.154			0.147		
Return Loss	-0.074			-0.069		
	(0.000)			(0.000)		

Note: P-values are in parentheses. This table summarizes portfolios that households manage themselves (Self-Managed) and portfolios they would receive if they become robo participants (Matched Robo), based on the Portfolio Dataset, which we use to calibrate the model in Section 6. Each observation is a pair of self-managed and robo portfolios. *Sharpe Ratio* is the ratio of expected return to standard deviation of return; *Expected Return* is the expected annual return based on a linear factor model, net of the risk-free rate; *Total Volatility* is the standard deviation of return; *Idiosyncratic Volatility* is the standard deviation of the pricing error in the factor model; and *Robo-Sharpe Expected Return* is the counterfactual expected return assuming that the self-managed portfolio invests with the same amount of total volatility as its robo match, which equals the product of the robo Sharpe ratio and the self-managed total volatility. *Return Loss* is the difference between the self-managed portfolio *Expected Return* and its *Robo-Sharpe Expected Return*. The baseline factor model, which is used in this table, is the Fama-French Three Factor Model augmented with U.S. and global bond returns (Fama-French with Bond). Columns (1)-(2) report the mean across portfolios for households in the second or third U.S. wealth quintiles, and columns (4)-(5) do so for the fourth and fifth quintiles. Columns (3) and (6) test for a difference in mean between matched robo and self-managed portfolios for each wealth class. The sample consists of non-advised portfolios for households who consult the robo advisor for a free portfolio review. Of these households, 45% become robo participants. Details on estimating the factor models are in Appendix D. The remaining notes are the same as in Table 1.

Table 7: Explaining the Data with the Model.

	Model	Data	Source
	(1)	(2)	(3)
(a) Baseline Model, Extensive Margin:			
Pre-Reduction Share from Middle-Class	14.9%	13.4%	Deposits Dataset
Post-Reduction Share from Middle-Class	43.1%	28.4%	
(b) Baseline Model, Intensive Margin:			
Robo Share, Existing Middle-Class	32.3%	31.1%	Deposits Dataset
Robo Share, New Middle-Class	21.4%	30.6%	
(c) Model with Participation Costs:			
Stock Market Participation, Middle-Class	17.3%	19.1%	SCF

Note: This table compares various model-implied statistics with their empirical counterparts, which assesses whether the quasi-experimental evidence can be rationalized using the portfolio choice model from Section 6. Panel (a) summarizes the share of robo participants from the middle class before and after the reduction, calculated as follows: in the model, we compute the share of participants from the middle class among those who find it optimal to participate under the previous minimum and under the reduced minimum; in the data, we compute the share of participants from the middle class among those who participate before the reduction in minimum and who become participants after the reduction. Panel (b) summarizes the average portfolio share allocated to the robo advisor for existing and new middle-class robo participants, calculated as follows: in the model, we compute the average robo portfolio share $\alpha_{i,t}^{A}$ among middle and upper class participants who find it optimal to participate under the reduced minimum; in the data, we compute the ratio of robo investment to liquid assets among middle class households during the pre-reduction period, for existing participants, and over the post-reduction period, for new participants. Panel (c) summarizes the share of middle-class households who participate in the stock market, calculated as follows: in the model, we compute the share of middle-class households who invest in either the robo portfolio or the self-managed portfolio before the reduction; in the data, we compute the share of middle-class households who invest in the stock market, based on the 2016 SCF as described in Appendix A. The statistics in panel (c) are based on a variant of the baseline model in which households incur a per-period cost of $\kappa = \$200$ when investing in the self-managed portfolio (Section C.4), since it is well-known that portfolio choice models generally do not deliver limited stock market participation without a fixed cost. All model-implied statistics aggregate household-level policy functions across the bins of age and liquid assets that define the state space, weighting by the share of households in the 2016 SCF within each bin. The remaining notes are the same as in Table 2.

Table 8: Welfare Implications of the Reduction

	Increase in Lifetime Consumption				
		Middle Class			Upper Class
	Pooled	Age 25-35	Age 36-55	Age 56-65	
	(1)	(2)	(3)	(4)	(5)
(a)Welfare Gain From the Reduction:	0.77%	0.58%	0.59%	1.68%	0%
(b) Gain from Reduction on Counterfactual Robo Portfolios:					
(i) Asset Allocation Counterfactuals					
Same Uncompensated Risk as Self-Managed Portfolio	0.23%	0.11%	0.10%	0.81%	
Same Market Loading as Self-Managed Portfolio	0.72%	0.55%	0.52%	1.62%	
Same Value and Size Loadings as Self-Managed Portfolio	0.46%	0.32%	0.27%	1.32%	
Same Bond Loadings as Self-Managed Portfolio	0.20%	0.15%	0.10%	0.58%	
(ii) Glide Path Counterfactuals					
No Adjustment for Age	0.37%	0.33%	0.23%	0.88%	
No Adjustment for Wealth	0.34%	0.34%	0.20%	0.73%	
No Adjustment for Age or Wealth	0.24%	0.26%	0.15%	0.45%	

Note: This table summarizes the average welfare gain for households who participate with the robo advisor under the reduced minimum but not under the previous minimum, based on the portfolio choice model in Section 6. Welfare gains are measured by the percent increase in annual consumption under the previous minimum that raises a household's expected lifetime utility by the same amount as the reduction, as in equation (20). Panel (a) summarizes the average of this statistic for all new middle-class participants in column (1), and columns (2)-(4) do so for various age bins. Panel (b) summarizes the welfare gain when the minimum is reduced on a counterfactual robo portfolio, which helps assess the channels that contribute to the result in (a). The counterfactual portfolios in panel (b.i) constrain the parameters of the robo portfolio for household i to equal those of the self-managed portfolio for household i. The constraints are: the same uncompensated risk ($\sigma^{\mathcal{A}}_{\epsilon,i} = \sigma^{\mathcal{S}}_{\epsilon,i}$); the same loading on the market factor ($\beta^{\mathcal{A}}_{MKT,i} = \beta^{\mathcal{S}}_{MKT,i}$); the same loading on value and size factors ($\beta^{\mathcal{A}}_{HML,i} = \beta^{\mathcal{S}}_{HML,i}$,$\beta^{\mathcal{A}}_{SMB,i} = \beta^{\mathcal{S}}_{SMB,i}$); and same the loadings on the U.S. and global bond factors ($\beta^{\mathcal{A}}_{USB,i} = \beta^{\mathcal{S}}_{USB,i}$,$\beta^{\mathcal{A}}_{GLB,i} = \beta^{\mathcal{S}}_{GLB,i}$). The counterfactual portfolios in panel (b.ii) shut down adjustment in the paramters of the robo portfolio ($\beta^{\mathcal{A}}_{i}, \sigma^{\mathcal{A}}_{\epsilon,i}$) by households characteristics. The adjustments that are shut down are: age, so that robo portfolios only adjust for liquid assets (No Adjustment for Age); liquid assets, so that robo portfolios only adjust for age (No Adjustment for Wealth); and both age and liquid assets, so that robo portfolios do not adjust at all (No Adjustment for Age or Wealth). The remaining notes are the same as in Table 2.

Table 9: Welfare Implications under Model Extensions

	Increase in Lifetime Consumption			
		Middle Class		
	Pooled	Age 25-35	Age 36-55	Age 56-65
	(1)	(2)	(3)	(4)
(a) Extensions				
Self-Managed Participation Cost				
GMM Cost ($\kappa = \$200$)	1.23%	0.89%	0.91%	2.87%
Higher Cost ($\kappa = \$700$)	1.24%	0.89%	0.91%	2.90%
Borrowing	0.80%	0.60%	0.60%	1.67%
Defined Contribution Plan	TBD	TBD	TBD	TBD
(b) Calibration of Portfolio Parameters				
Heckman (1979) Selection Model	0.92%	0.67%	0.68%	2.06%
Calvet, Campbell and Sodini (2007)	0.88%	0.64%	0.66%	1.97%
(c) Calibration of Other Parameters				
GMM Estimates ($\gamma = 5.1$, $\delta = 0.92$)	1.19%	0.84%	0.97%	2.38%
High Impatience ($\delta = 0.85$)	0.74%	0.56%	0.59%	1.53%
High Labor Income Loading ($\beta^Y = 0.13$)	0.77%	0.58%	0.59%	1.68%
(d) Gain from Comparison Shocks				
+1.5 pps in Equity Premium	0.82%	0.49%	0.89%	1.33%
+5 pps in Log Permanent Income	0.93%	0.58%	0.88%	1.97%

Note: This table recalculates the welfare gains from the reduction under various extensions and reparameterizations of the baseline model in Table 8. Panel (a) summarizes the average welfare gain under models with the following extensions: a per-period cost of κ when holding the self-managed portfolio (Self-Managed Participation Costs); the ability to borrow up to 30% of one's liquid assets at the average rate on credit card debt in 2015 of 12% (Borrowing); and, following Campbell et al. (2001), a requirement to allocate 10% of one's income to a defined contribution plan that delivers the same annual return as the risk-free asset and cannot be liquidated until retirement (Defined Contribution). The participation cost κ equals either \$200, based on the GMM estimates in Appendix Table A9, or a higher value of \$700. Panel (b) assesses robustness to the calibration of the self-managed portfolio parameters β_i^S and $\sigma_{\epsilon,i}$ in two ways, both of which modify the projection of portfolio characteristics onto the model's state variables described in Section 6.2.1. First, we adjust for selection into knowing about the robo advisor by controlling for the inverse Mills ratio in this projection, which comes from the Heckman (1979) procedure described in Section 5.4. Second, we use the projection summarized in Table 5 of Calvet, Campbell and Sodini (2007) to calibrate a self-managed portfolio's loading on the market factor and its idiosyncratic volatility as a function of the model's state variables. Panel (c) assesses robustness to the calibration of other parameters. Note that a loading of log labor income of financial returns of $\beta^Y = 0.13$ corresponds to a correlation coefficient of 20%. To assess the magnitude of the results in Table 8, panel (d) summarizes the welfare gain from comparison shocks that occur under the previous minimum: a permanent 1.5 pps increase in the expected excess return on the U.S. stock market; and a 5 pps increase in permanent labor income, equal to one-half the standard deviation of permanent labor income shocks ($\nu_{i,t}$). The remaining notes are the same as in Table 8.

Online Appendix

This document contains additional material referenced in the text. Appendix A describes our data in greater detail. Appendix B contains details on our core microeconometric analysis from Section 4 and performs the robustness exercises referenced in Section 5. Appendix C provide additional information about the life-cycle model from Section 6. Appendix D describes the method for estimating the idiosyncratic volatilities and factor loadings on self-managed and robo portfolios in Section 6.2.1. Additional figures and tables are at the end of the appendix.

A Data Appendix

We provide additional details on the paper's two principal datasets: a weekly panel of deposit activity by robo participants (A.1); and a cross-section of portfolio snapshots for self-managed, non-robo portfolios (A.2). We also describe other datasets (A.3) and provide a catalog of the paper's key variables (A.4).

A.1 Deposits Dataset

Our first robo advising dataset contains a weekly time series of deposits with the robo advisor, Wealthfront. We obtain this information directly through a query of Wealthfront's internal server. The query merges two internal subdatasets. The first subdataset includes demographic information about Wealthfront participants. The second subdataset contains the date and size of each deposit made by a Wealthfront participant from December 1, 2014 through February 29, 2016. The internal query then merges these two subdatasets together based on username and tax status of the portfolio associated with the username. Each merged observation defines a "robo participant". As implied by Table 1, the merged dataset includes information on 9,702 Wealthfront participants who made at least one deposit during the sample period, 4,366 of whom became participants before the July 2015 reduction and 5,336 of whom become participants afterward.

Summarizing the discussion in the text, we observe the date and size of the deposit and whether the deposit comes from a new participant. In addition, we observe the participating household's annual income, state of residence, liquid assets, recommended and selected risk tolerance score, and householder age. Per the language of the questionnaire, liquid assets are defined as "cash, savings accounts, certificates of deposit, mutual funds, IRAs, 401ks, and public stocks".

The risk tolerance score defines the portfolio allocation received by the participant, as shown in Table A2. The recommended risk tolerance score is a function of the household's demographic information and answers to several questions about financial goals and response to market downturns. The selected risk tolerance score equals the recommended score for 64% of Wealthfront participants, and the remaining participants select a different score. We use this difference to calculate a measure of high risk aversion, denoted $Risk\ Averse_i$ in the text. Only 3% of households who select a different risk tolerance score deviate from their recommended score by more than 3 points, corresponding to a shift in CAPM beta of around 15 pps.

We cross-referenced our robo advising dataset against publicly available SEC ADV filings. According to these filings, Wealthfront reported 18,800 participants (i.e., clients) in December 2014 and 61,000 participants in February 2016. As described in the text, the discrepancy between the SEC ADV filings and our dataset is explained by the SEC's filing requirements. Specifically, the SEC states: "The definition of 'client' for Form ADV states that advisors must count clients who do not compensate the advisor" (SEC 2017). Thus, the number of participants reported to the SEC by Wealthfront or any other robo advisor

includes participants who did not make any deposits over the sample period as well as "participants" who created a username but never funded a Wealthfront account.

A.2 Portfolio Dataset

Our second robo advising dataset covers a subset of households in the deposits dataset as well as a set of robo non-participants. It contains snapshots of households' portfolio holdings in an outside, traditional brokerage account. This information is paired with the portfolio holdings of the household's counterfactual robo portfolio, along with the same demographic information as in the deposits dataset. We also observe each portfolio's advisory fees and tax status. The dataset was generated by a free online tool through which our data provider gave financial advice to candidate clients about their outside portfolio holdings.

Specifically, candidate clients would provide their log-in credentials for their outside brokerage 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, our data provider would ask the household to answer its standard questionnaire, which is the source of our demographic variables. Finally, at the conclusion of the report, our data provider would tell the household the portfolio she would receive as a client and give her the option to fund a robo portfolio. This produces a matched, counterfactual robo portfolio for each household in the sample. Our sample contains 2,654 snapshots taken within a window of the reduction in account minimum.

Given that we use the dataset to calibrate self-managed portfolio characteristics, we filter this dataset to only include pairs of portfolios in which the non-robo portfolio has no advisory fee. The filtered dataset includes 1,913 portfolios, as shown in Table 6.

A.3 Auxiliary Datasets

A.3.1 Survey of Consumer Finances

The Survey of Consumer Finances (SCF) is a publicly available dataset administered by the Federal Reserve Board every three years, and we rely on the 2016 dataset. The SCF contains financial and demographic information about a representative cross-section of U.S. households. The SCF is one of the most commonly used datasets in the literature, and Bricker et al. (2017) provide a thorough overview of it.

We use the SCF dataset to calculate quintiles of the overall U.S. distribution of liquid assets. To maximize comparability with our robo advising dataset, we define liquid assets in the SCF as the sum of checking accounts, savings accounts, certificates of deposit, cash, stocks, bonds, savings bonds, mutual funds, annuities, trusts, IRAs, and employer-provided retirement plans. This definition of liquid assets most closely matches the definition in our robo advising dataset, although the two are not equivalent. For example, we include bonds and savings bonds in the SCF definition, although they are not explicitly mentioned as a liquid asset in the robo advisor's questionnaire. Removing bonds and savings bonds from the SCF definition has little impact because it only changes the boundary between the middle and upper classes by 1%. We carefully examine how measurement error might affect our results in Section 4.3. Appendix Table A3 reports the boundaries that define the five quintiles.

A.3.2 Asset Pricing Datasets

We use data from the CRSP annually updated stock file, Ken French's website, and the Bloomberg-Barclays aggregate U.S. and unhedged global bond indices to estimate the asset pricing factor models, as described in Appendix D.

A.3.3 Other Auxiliary Datasets

We use the Center for Disease Control's mortality tables to calibrate the survival probabilites in the portfolio choice model (Xu et al. (2020)). The CDC reports these probabilities for brackets of the age distribution, and we use the average within each bracket. We calculate the survival probability as one minus the mortality rate. For the post-retirement period (i.e., $t \geq \underline{T}$), we use the lowest survival probability across age brackets. Lastly, we calibrate the credit card interest rate of 12% using the 2015 Federal Reserve Consumer Credit Report.

A.4 Description of Variables

Our empirical analysis relies on the following variables:

- *Liquid Assets*$_i$: This variable is the sum of cash, savings accounts, certificates of deposit, mutual funds, IRAs, 401ks, and public stocks for household i, based on the deposits dataset. We interpret "401k" to also mean other tax-advantaged, employer-sponsored plans, such as 403b or 457 plans.
- *Middle*$_i$: This variable indicates if household i's liquid assets fall within the second or third U.S. quintile of liquid assets. Household i's liquid assets are calculated using the deposits dataset. Quintiles of liquid assets are calculated using the SCF dataset.
- *New Participant*$_i$: This variable indicates if household i becomes a participant with the robo advisor over the period from July 7, 2015 through February 29, 2016. Explicitly, it equals 1 for such households and equals 0 for households who participated before July 7, 2015.
- *Initial Deposit*$_i$: This variable is the initial deposit with the robo advisor made by household i, based on the deposits dataset.
- *Income*$_i$: This variable is annual income for household i, based on the deposits dataset.
- *Age*$_i$: This variable is the age of the householder for household i, based on the deposits dataset.
- *Risk Averse*$_i$: This variable indicates if household i chooses a lower risk tolerance score than recommended by the robo advisor, based on the deposits dataset.
- *Under Minimum*$_i$: This variable indicates if household i's initial deposit with the robo advisor is less than \$5,000, based on the deposits dataset.
- *At Minimum*$_i$: This variable indicates if household i's initial deposit with the robo advisor is between \$5,000 and \$5,250, based on the deposits dataset.

B Econometric Appendix

We provide the details for aggregate results presented in Table A5 (B.2). Then we explain how the estimation of equation (2) can be affected by measurement error (B.4). Lastly, we discuss the robustness of our main results from Section 4 to the effects of business stealing (B.1) and gambling motives (B.5).

B.1 Business Stealing: Reallocation of Participants across Advisors

New middle-class robo participants may have planned to invest with a competitor robo advisor during the post-reduction period, but the reduction prompted them to invest with Wealthfront instead. In this case, our results would reflect business stealing rather than democratization of automated wealth management. We assess this possibility by using data from the SEC's Form ADV to plot new participants at other standalone robo advisors, namely Betterment and Personal Capital. While the Form ADV data have limitations described in Section 3, they are the best source of data for this exercise, short of having microdata from each major U.S. robo advisor. The results in Appendix Figure D show very little change in trend growth in the number of participants at Wealthfront's competitors from the pre-reduction to the post-reduction periods. This observation suggests that the reduction indeed expands access to automated wealth management, rather than simply reallocating participants across robo advisors.

B.2 Details on Aggregation Exercise to Interpret the Magnitude

We provide details on the aggregation exercise referenced in Section 5.3, which is intended to help interpret the magnitude of the baseline results in Table 2. We would like to decompose the observed growth rate in the total number of robo participants into the component due to the reduction versus that due to other forces. Let g denote the observed growth rate, which we calculate directly from the data. Consider a counterfactual without the reduction, in which middle-class households do not experience a relaxation of minimum-account constraints and, thus, $\mu = 0$. Note that the observe growth rate in the total number of robo participants can be directly calculated from the data as

$$g = \frac{\textit{New Participants}}{\textit{Existing Participants}}, \tag{B1}$$

where *New Participants* is the number of households who become robo participants after the reduction; and, analogously, *Existing Participants* is the number who participated beforehand. It will be helpful to rewrite the numerator of equation (B1) as

$$\textit{New Participants} = \mathbb{E}\left[\textit{New Participant}_i\right] \times \textit{All Participants}, \tag{B2}$$

where $\textit{All Participants} = \textit{New Participants} + \textit{Existing Participants}$ is the sum of new and existing robo participants; and $\mathbb{E}\left[\textit{New Participant}_i\right]$ is the share of all such robo participants who are new. Substituting equations (B2) and (2) into equation (B1) allows us to express g as

$$g = \frac{\mathbb{E}\left[\textit{New Participant}_i\right]}{1 - \mathbb{E}\left[\textit{New Participant}_i\right]} = \frac{\mu\mathbb{E}\left[\textit{Middle}_i\right] + \psi\mathbb{E}\left[X_i\right] + \varrho}{1 - \left(\mu\mathbb{E}\left[\textit{Middle}_i\right] + \psi\mathbb{E}\left[X_i\right] + \varrho\right)}, \tag{B3}$$

which, by definition, is numerically equivalent to the expression in equation (B1).

Consider a counterfactual without the reduction, in which middle-class households do not experience a relaxation of constraints and, thus, $\mu = 0$. Under this counterfactual, the overall number of robo

participants grows at the rate

$$g^C = \frac{\psi \mathbb{E}[X_i] + \varrho}{1 - (\psi \mathbb{E}[X_i] + \varrho)}, \tag{B4}$$

or, equivalently,

$$g^C = \frac{\mathbb{E}[New\ Participant_i] - \mu \mathbb{E}[Middle_i]}{1 - (\mathbb{E}[New\ Participant_i] - \mu \mathbb{E}[Middle_i])}. \tag{B5}$$

When restricting the focus to the number of middle-class participants, we have similar expressions

$$g^C\Big|_{Middle} = \frac{\psi \mathbb{E}[X_i | Middle_i = 1] + \varrho}{1 - (\psi \mathbb{E}[X_i | Middle_i = 1] + \varrho)}, \tag{B6}$$

which we can rewrite as

$$g^C\Big|_{Middle} = \frac{\mathbb{E}[New\ Participant_i | Middle_i = 1] - \mu}{1 - (\mathbb{E}[New\ Participant_i | Middle_i = 1] - \mu)}.$$

Our statistic of interest is

$$\eta \equiv g - g^C, \tag{B7}$$

which equals the component of the observed growth in the total number of robo participants that is due to the reduction.

Appendix Table A5 summarizes various calculations of η and of the analogous statistic for growth in middle-class households' robo participation. Interpreting the first row, the baseline estimates from Table 2 imply that the reduction increases the overall number of robo participants by 13%, which is driven by a 108% increase in the number of middle-class participants.[17] The additional estimates from Table 2 imply an increase in the number of middle-class participants between 127% and 148%. The economic discussion of these results is in the text in Section 5.3.

B.3 Details on External Validity Exercise

We provide details for the exercise mentioned in Section 5.4, which implements a Heckman (1979) selection estimator to assess how the results might apply to a broader population that does not endogenously know about the robo advisor. To structure the argument, suppose that a household knows about the robo advisor if some latent variable, $Aware_i^*$, exceeds zero. This variable evolves similarly to equation (1),

$$Aware_i^* = \iota_0 + \iota_1 Middle_i + \iota_2 X_i + z_i. \tag{B8}$$

where z_i is an unobserved characteristic of those who know about the robo advisor. This characteristic is normally distributed and potentially correlated with unobserved determinants of participating after the reduction, u_i. In particular,

$$u_i = \varphi z_i + \tilde{u}_i, \tag{B9}$$

where $\tilde{u}_i$ is an idiosyncratic participation shock.

[17] In relation to Table 2, the 108% increase in the number of middle-class participants follows from the estimated 14 pps increase in their probability of participation because the middle class was underrepresented before the reduction.

Thus, the households in our data possess a sufficiently high value of z_i that $Aware_i^* > 0$. Heckman (1979) shows how the conditional mean of z_i for such households equals $\lambda(\iota_0 + \iota_1 Middle_i + \iota_2 X_i)$, where $\lambda(a)$ denotes the inverse Mills ratio evaluated at a. Therefore, reestimating equation (1) after controlling for $\lambda(\iota_0 + \iota_1 Middle_i + \iota_2 X_i)$ allows us to identify μ in the presence of sample selection. Intuitively, the propensity to know about the robo advisor, captured by the inverse Mills ratio, is an omitted variable that may affect whether a household participates after the reduction.

We implement this procedure as follows. First, we use the SCF dataset to estimate a probit regression that predicts whether a household seeks financial advice online. This variable is the best proxy for whether a household knows about the robo advisor (i.e., $Aware_i^* > 0$), among variables in the SCF dataset. The predictor variables are those in Table 2 that also appear in the SCF, age and log income. Next, we calculate the inverse Mills ratio for each household in our robo advisory dataset and use it as a control variable when reestimating the baseline specification from Table 2. We bootstrap the entire procedure to obtain standard errors.

Appendix Table A10 summarizes the results. The estimate of μ falls by 3 pps to 12.5%, as mentioned in the text.

B.4 Measurement Error

As mentioned in the text, the variable $Middle_i$ may be subject to additive measurement error due to self-reporting. On the one hand, such measurement error introduces attenuation bias, which would tend to bias the estimates toward zero. Similarly, the estimates are biased toward zero if new robo participants overreport their wealth more than existing participants do. On the other hand, measurement error biases the estimates away from zero if new participants underreport their wealth relative to existing participants. Formally, if we mis-measure $Middle_i$ as $\widehat{Middle}_i = Middle_i + \varepsilon_i$, then the estimator for μ in a specification of equation (2) without controls is

$$\widehat{\mu} = \mu \left(1 - \frac{\text{Var}\left[\varepsilon_i\right] + \mathbb{E}\left[Middle_i \times \epsilon_i\right]}{\text{Var}\left[\widehat{Middle}_i\right]} \right) + \frac{\mathbb{E}\left[u_i \times \varepsilon_i\right]}{\text{Var}\left[\widehat{Middle}_i\right]}. \tag{B10}$$

The term in parentheses captures the effect of attenuation bias. The second term captures bias from differences in misreporting between new and existing participants.

B.5 Gambling

Experimental evidence suggests that households exhibit lower risk aversion in the context of small lotteries (e.g., Bombardini and Trebbi 2012). Therefore, the baseline results are unlikely to confound gambling motives, since such motives would be stronger among upper-class households, for whom an investment under $5,000 is relatively small. If anything, such a gambling channel would imply negative estimates, which is not in line with the results.

C Model Appendix

We describe the solution of the model in Section 6 (C.1), derive the welfare measure studied in Section 8 (C.3), provide details on the model extensions also studied in Section 8 (C.4), and structurally estimate the model's preference parameters as referenced in Section 6.2 (C.5).

C.1 Model Solution

We first restate the problem, and then we describe the numerical solution algorithm.

C.1.1 Consolidated Problem

Repeating from Section 6.1.4, households i of age t solves the following Bellman equation,

$$V_t(W_{i,t}) = \max_{\alpha^f_{i,t},\alpha^S_{i,t},\alpha^A_{i,t}} \left\{ \frac{{C_{i,t}}^{1-\gamma}}{1-\gamma} + \delta p_t \mathbb{E}_t \left[V_{t+1}(W_{i,t+1}) \right] \right\} \tag{C1}$$
$$s.t. (8)\text{-}(16)$$

for $t_0 \leq t \leq \underline{T}$. Recall that households solve an eat-the-pie income for $t > \underline{T}$. Explicitly, they solve

$$V_{\underline{T}+1}(W_{i,\underline{T}+1}) = \max_{\left\{\alpha^f_\tau\right\}} \sum_{\tau=\underline{T}+1}^{\tau=\overline{T}} (\delta p_{\underline{T}+1})^{\tau-\underline{T}-1} \frac{{C_{i,\tau}}^{1-\gamma}}{1-\gamma} \tag{C2}$$
$$s.t.$$
$$0 \leq \alpha^f_\tau \leq 1 \tag{C3}$$
$$W_{i,\tau+1} = \alpha^f_{i,\tau}(1+R^f)W_{i,\tau} \tag{C4}$$
$$C_{i,\tau} = [1-\alpha^f_{i,\tau}]W_{i,\tau}. \tag{C5}$$

Indirect utility has the familiar form

$$V_{\underline{T}+1}(W_{i,\underline{T}+1}) = B\frac{{W_{i,\underline{T}+1}}^{1-\gamma}}{1-\gamma}, \tag{C6}$$

with

$$B = \sum_{\tau=\underline{T}+1}^{\overline{T}} \delta^{\tau-\underline{T}-1} \left[\frac{1-\chi}{1-\chi^{\overline{T}-\underline{T}}}\right] \left[\delta(1+R^f)\right]^{\frac{\tau-\underline{T}-1}{\gamma}}, \tag{C7}$$
$$\chi = \delta^{\frac{1}{\gamma}}(1+R^f)^{\frac{1-\gamma}{\gamma}} \tag{C8}$$

as, for example, shown in Costa-Dias and O'Dea (2019).

C.1.2 Numerical Algorithm

Our numerical algorithm is standard and follows the methods typically used in benchmark portfolio choice models. First, we solve equation (C1) for age $\underline{T}$ as a function of liquid assets: $V_{\underline{T}}(W_{i,\underline{T}})$. Since the

solution does not have an analytic form, we discretize liquid assets. The grid ranges from the minimum value of liquid assets in the 2016 SCF to the 90$^{\text{th}}$percentile of liquid assets in increments of 0.1 on a log scale. The resulting grid has 109 points. This discretization intentionally places most of its density in the bottom four quintiles. Our empirical results imply that the strongest response to the reduction will occur in this region, and so we want to minimize approximation error in it. We obtain very similar theoretical results under alternative discretizations. We also follow convention by discretizing the set of shocks and approximating their joint distribution through Gaussian quadrature (e.g., Tauchen and Hussey (1991)). For completeness, the model's shocks are: $F_t, \epsilon^{\mathcal{S}}_{i,t}, \epsilon^{\mathcal{A}}_{i,t}$, and $\overline{\nu}_{i,t} \equiv \xi_{i,t} + \nu_{i,t}$. The grid has 5 points for each shock.

Following standard practice, we optimize by grid search, and so we avoid selecting local optima. Accordingly, we discretize each of the control variables, $\alpha^{f}_{i,t}, \alpha^{\mathcal{S}}_{i,t}$, and $\alpha^{\mathcal{A}}_{i,t}$, into a 40 point grid. We can further reduce the dimensionality of the resulting 40×3 matrix by omitting choices that violate one of the constraints (8)-(16). As mentioned in the text, we simplify the model's computational complexity by assuming households cannot hold the self-managed and automated portfolios concurrently: $\min\left\{\alpha^{\mathcal{S}}_{i,t}, \alpha^{\mathcal{A}}_{i,t}\right\} = 0$. We obtain similar results without this simplification because it is rarely optimal to hold both at the same time. We also reduce computational complexity by assuming households must maintain a minimum balance of M with the automated asset manager, whereas, in reality, households only need to make an initial deposit of M. Otherwise, we must keep track of $\alpha^{\mathcal{A}}_{i,t}$as an additional state variable because it determines the lower bound on a household's robo investment. Namely, under an initial deposit requirement, households do not need to top-up their balance to M if market forces push their account balance below this threshold. We assess the validity of this simplification by solving the model under the more realistic yet intensive setup with a minimum deposit requirement, finding similar results as in Table 8. Intuitively, the high expected return on the robo portfolio makes cases of a top-up relatively rare.

Next, after solving $V_{\underline{T}}(W_{i,\underline{T}})$, we iterate backward, solving $V_{\underline{T}-1}(W_{i,\underline{T}-1})$ and so forth until we arrive at the initial period, t_0. For each age t, we approximate $V_{t+1}(W_{i,t+1})$ using a linear spline interpolation in liquid assets, $W_{i,t+1}$, and we evaluate $\mathbb{E}[V_{t+1}(W_{i,t+1})]$ using Gaussian quadrature, as mentioned above. The presence of a binding account minimum in the middle of the state space complicates our problem, but the linear spline handles this challenge well, producing well-behaved policy functions. By contrast, a cubic spline generates policy functions that are less well-behaved. Moreover, a linear spline can still generate substantial curvature in the utility function for a suitably fine discretization of the state space. Lastly, we solve $V_t(W_{i,t})$ using the labor income parameters shown in Table 5, setting income equal to the median income in the 2016 SCF for the baseline cohort studied in Cocco, Gomes and Maenhout (2005).

Summarizing, this algorithm results in a sequence of value functions $\{V_\tau(W_{i,\tau})\}$ and policy rules $\left\{\alpha^{f}_{\tau}(W_{i,\tau}), \alpha^{\mathcal{S}}_{\tau}(W_{i,\tau}), \alpha^{\mathcal{A}}_{\tau}(W_{i,\tau})\right\}$ that we use in the positive and welfare analyses of Sections 7 and 8.

C.2 Discussion of Policy Functions and Life Cycle Intuition

Appendix Figure D summarizes the model's policy functions for the following outcomes: decision to participate with the robo advisor (Robo Participation); share of investible wealth allocated to the robo portfolio (Robo Share); share of investible wealth allocated to either the self-managed portolfio or the robo portfolio (Risky Share); and share of liquid assets that is saved (Savings Rate), recalling that we use "liquid assets" and "cash-on-hand" synonymously. By "investible wealth", we mean the difference between liquid assets and optimal consumption. We report the optimal choice of each outcome for pairs of age and liquid assets, which are, effectively, the model's two state variables. We consider policy functions before the reduction, in which the account minimum equals \$5,000.

First, consider robo participation in the upper-left panel. We plot the minimum threshold of liquid assets at which the household decides to participate with the robo advisor as a function of household

age. This threshold remains relatively stable at around $27,000 until the household reaches peak earning years. Notably, the fact that the threshold exceeds the account minimum by a factor of 5 accords with the empirical observation that the reduction has a large effect on households from the third quintile of the U.S. wealth distribution, which is bounded by $6,000 and $42,000. After age 40, the threshold starts to decline, meaning that older households are willing to invest a larger share of their liquid assets with the robo advisor. This reflects how households' consumption-to-wealth ratio falls as they approach retirement, or, equivalently, their savings rate increases, which is a well-known life cycle pattern generated by benchmark models summarized by Gomes (2020). Our model also generates this pattern, shown in the lower-right panel of Appendix Figure D.

In benchmark models, households reduce their risky share as they age to offset a decline in bond-like labor income. According to canonical life cycle intuition, this channel should increase the robo participation threshold as households age, thus working in the opposite direction as the savings rate channel described in the previous paragraph. However, robo portfolios differ from the risky asset featured in most life cycle models, the latter of which is modelled as a stock index fund. In particular, robo portfolios feature a glide path towards less total risk as households age. For example, recall how robo porfolios for older or less-wealth households have a stronger tilt towards bonds and away from stocks (Figure 4). This glide path feature incentivizes older households to invest a larger share of their wealth in the robo portfolio. In fact, the upper-right panel of Appendix Figure D implies that, beginning around age 40, the optimal robo share increases as households age, holding liquid assets fixed. This increase in robo share drives an increase in risky share, as we now describe.

The lower-left panel of Appendix Figure D plots the optimal risky share as a function of age and wealth. Consistent with benchmark models, the optimal risky share declines in age until the household reaches peak earning years, which is best seen by the lightening of color as one moves up the vertical axis holding liquid assets fixed at around $25,000. In addition, the optimal risky share weakly increases in wealth. Since we do not incorporate fixed costs of stock market participation until Section C.4, the optimal risky share is always positive, even for very low levels of wealth. Note, however, that most households in the middle class select a risky share well below 100%. Benchmark models, by contrast, predict a 100% risky share for almost all households. This difference arises because households in benchmark models invest in the total stock market, and, thus, have a relatively low level of idiosyncratic risk. By contrast, less-wealthy households in our setting must invest in the self-managed portfolio, which, per Table 6, has a relatively low Sharpe ratio. This underdiversification reduces their demand for risky assets. However, once liquid assets reaches around $27,000, households begin to invest in the robo portoflio, and the the risky share jumps to 100%.

Lastly, we point out three features of the savings rate policy function in the lower-right panel that accord with benchmark life cycle models. First, as already described, the savings rate begins to increase in age after the household reaches her peak earning years, holding liquid assets fixed. Second, the savings rate decreases in age leading up to peak earning years, again holding liquid assets fixed. This pattern reflects how the household wishes to accumulate a buffer against idiosyncratic labor income shocks. Such precautionary savings arises because isoelastic preferences feature prudence, as controlled by γ. Third, the savings rate increases in liquid assets, which is another standard feature of benchmark models.

Summarizing, the model's policy functions resemble those of benchmark models, with several exceptions that are due to our unique setting. In particular, the combination of an account minimum and inefficiency in self-managed portfolios generates a jump in the risky share once the household's liquid assets reach some threshold. Benchmark models feature no such jump. Moreover, a personalized risk profile by both age and wealth increases older households' and less-wealthy households' demand for risky assets beyond what a benchmark model would predict.

C.3 Welfare Measure

Repeating from the text, we measure household i's welfare gain by the percent increase in annual consumption under the previous minimum that raises her expected lifetime utility by the same amount as the reduction, denoted by q_i . Let $\{\underline{C}_{\{i,\tau\}}\}_{\tau\geq t}$ denote the optimal consumption stream for household i under the previous minimum of \$5,000 and let $\{\overline{C}_{\{i,\tau\}}\}_{\tau\geq t}$ denote the optimal consumption stream under the reduced minimum of \$500. Then q_i is defined by solving

$$\mathbb{E}_t\left[\sum_{\tau=t}^{\tau=\overline{T}}\delta^{\tau-t}\left(\prod_{j=t}^{j=\tau-1}p_j\right)\frac{\left((1+q_i)\underline{C}_{i,\tau}\right)^{1-\gamma}}{1-\gamma}\right]=\mathbb{E}_t\left[\sum_{\tau=t}^{\tau=\overline{T}}\delta^{\tau-t}\left(\prod_{j=t}^{j=\tau-1}p_j\right)\frac{\left(\overline{C}_{i,\tau}\right)^{1-\gamma}}{1-\gamma}\right]\equiv\overline{V}_i, \tag{C9}$$

where, as in the text, $\overline{V}_i$ denotes household i's value function under the \$500 minimum. Likewise, let

$$\underline{V}_i\equiv\mathbb{E}_t\left[\sum_{\tau=t}^{\tau=\overline{T}}\delta^{\tau-t}\left(\prod_{j=t}^{j=\tau-1}p_j\right)\frac{\left(\underline{C}_{i,\tau}\right)^{1-\gamma}}{1-\gamma}\right] \tag{C10}$$

denote i's value function under the \$5,000 minimum. Rearranging terms gives

$$q_i=\left(\frac{\overline{V}_i}{\underline{V}_i}\right)^{\frac{1}{1-\gamma}}-1, \tag{C11}$$

as shown in equation (20). Note that equation (C11) is increasing in the difference between $\overline{V}_i$ and $\underline{V}_i$because a standard isoelastic utility function is bounded above by zero.

C.4 Model Extensions

Section 8.1 considers three model extensions that we now describe. Panel (c) of Table 8 summarizes welfare gains under these extensions.

C.4.1 Participation Costs

The Participation Cost extension introduces a per-period utility loss from self-management equal to a $\kappa = \$100$ reduction in consumption. Explicitly, flow utility becomes: $\frac{(C_{i,t}-\kappa\mathbf{1}\cdot[\alpha^{\mathcal{S}}_{i,t}>0])^{1-\gamma}}{1-\gamma}$. For reference, Vissing-Jørgensen (2003) finds a per-period cost of \$830 in 2015 dollars (\$350 in 1982-1984 dollars). We intentionally choose a low participation cost because our model already explicitly incorporates an account minimum and underdiversification. The remaining participation cost may be interpreted as the time cost associated with rebalancing.

C.4.2 Borrowing

The Borrowing extension allows households to borrow up to $b = 30\%$ of liquid wealth. Explicitly, constraint (13) becomes: $\alpha^f_{i,t}\geq -b$. One can interpret the borrowing limit, b, as reflecting how, for example, lender-friendly bankruptcy laws allow recourse up to 30% of liquid wealth. The interest rate at which households borrow equals 12%, the average interest rate on credit card debt in 2015, or 12%.

C.4.3 Defined Contribution Plan

The Defined Contribution Plan extension requires households to allocate 10% of their income to a plan that delivers the same annual return as the risk-free asset, R^f. This setup and parameterization follows

Campbell et al. (2001). Liquid wealth for household i of age t is defined to include the value of plan assets accumulated to that point. As in Campbell et al. (2001), households cannot withdraw funds from the plan until age $\underline{T}+1$.

C.5 Structural Parameter Estimation

We assess the validity of the preference parameter values in Table 5 by estimating these parameters through a generalized method of moments estimator (GMM). The 3×1 parameter vector is (γ, δ, κ). The 9×1 moment vector consists of: a 4×1 vector containing the share of robo participants from the second through the fifth U.S. wealth quintiles under the previous account minimum; an analogous 4×1 vector containing the shares under the reduced minimum; and the coefficient μ from an uncontrolled specification of equation (2), our baseline regression equation.[18] Let θ_m denote moment m, and let

$$\Theta = [\theta_1, ..., \theta_9]'$$

denote the full moment vector.

We solve for the value of (γ, δ, κ) that minimizes the weighted squared distance between the theoretical moment vector, $\tilde{\Theta}(\gamma, \delta, \kappa)$, and the empirical moment vector, $\hat{\Theta}$, using a similar solution technique as in Appendix C.1. Denote the solution by

$$(\hat{\gamma}, \hat{\delta}, \hat{\kappa}) = \arg \min_{(\gamma,\delta,\kappa)} (\tilde{\Theta}(\gamma, \delta, \kappa) - \hat{\Theta})' \Psi (\tilde{\Theta}(\gamma, \delta, \kappa) - \hat{\Theta}), \tag{C12}$$

where Ψ is a 9×9 weight matrix. Our baseline weight matrix is the GMM optimal matrix (Newey (2007)), in which moments are weighted by the inverse standard error of the empirical moment. Accordingly, we estimate $\hat{\gamma} = 5.1$, $\hat{\delta} = 0.92$, and $\hat{\kappa} = 169$ as shown in column (1) of Appendix Table A9. We obtain estimates of $\hat{\gamma} = 7$, $\hat{\delta} = 0.97$, and $\hat{\kappa} = 206$ when weighting moments equally, shown in column (2).

Column (1) estimates the parameters when weighting each empirical moment by its inverse standard error (GMM Optimal). For this reason, the share of participants from the first U.S. wealth quintile is excluded from the moment vector, as this share has no empirical variance. Column (2) estimates the parameters when equally weighting each empirical moment. Standard errors are bootstrapped. Brackets correspond to 95% confidence intervals under an asymptotically normal distribution.

[18] We do not include moments related to the share of robo participants from the first U.S. wealth quintile, since this share has no empirical variance and, thus, would have infinite weight in the GMM optimal weight matrix.

D Asset Pricing Appendix

This appendix describes the method for estimating the idiosyncratic volatilities and factor loadings on self-managed and robo portfolios in Section 6.2.1. We follow Calvet et al. (2007) in our methodology. For each security k in the portfolio dataset described in Section 3.2 and a given asset pricing model F, we estimate the following equation

$$R_{k,m} = \beta_k^F F_m + \epsilon_{k,m}^F, \tag{D1}$$

where F_m denotes a column vector of pricing factors in month m; β_k^F denotes the respective row vector of loadings; $R_{k,m}$ denotes the monthly return on security k in excess of the risk-free return, measured by the one-month Treasury yield, and net of expense ratios and other fees; and $\epsilon_{k,m}^F$ is an idiosyncratic, zero-mean shock to security k with standard deviation $\sigma_{\epsilon,k}^F$. We estimate equation (D1) using the longest available time series of monthly returns for each security k and factor vector F. Following Calvet et al. (2007), we restrict the set of securities to stocks and funds traded on an exchange, which has little material effect on the results because few portfolios have bonds, derivatives, or other less-liquid securities. Given the estimated loadings $\hat{\beta}_k^F$ from estimating equation (D1) for model F, it is straightforward to compute the idiosyncratic volatility and factor loadings for household i's self-managed and robo portfolios, as in equation (19).

We estimate equation (D1) for the following asset pricing models,

$$F_m^{CAPM} = \left[R_m^M\right]', \tag{D2}$$

$$F_m^{FF} = \left[R_m^M, \quad R_m^{HML}, \quad R_m^{SMB}\right]', \tag{D3}$$

$$F_m^{FF+} = \left[R_m^M, \quad R_m^{HML}, \quad R_m^{SMB}, \quad R_m^{USB}, \quad R_m^{GLB}\right]', \tag{D4}$$

where R_m^M is the monthly market return based on the global Morgan Stanley Capital International Index (MSCII), net of the one-month Treasury yield; R_m^{HML} is the spread in monthly return between high book-to-market stocks and low book-to-market stocks; R_m^{SMB} denotes the spread in monthly return between stocks with a small market capitalization and a big market capitalization; R_m^{USB} is the monthly return on the Barclays Aggregate U.S. Bond Index Unhedged, net of the one-month Treasury yield; and R_m^{GLB} is the monthly return on the Barclays Global Aggregate Bond Index Unhedged, net of the one-month Treasury yield.

Equations (D2)-(D4) are: the standard capital asset pricing model (CAPM), the Fama-French Three Factor Model, and a five-factor model augmenting the Fama-French model with U.S. and global bond returns. Our data on monthly returns come from the Center for Research in Security Prices (CRSP) and Ken French's website, as described in Appendix A. We use the sample mean to calibrate the factor risk prices, multiplied by 12 to obtain an approximate annual value. Analogously, we use the sample covariance matrix to calibrate the covariance of the factor vector. The moments of the factor vectors in equations (D2)-(D4) are shown in Appendix Table A7.

Additional Figures and Tables

Figure A1: Growth in Robo Participation by U.S. State

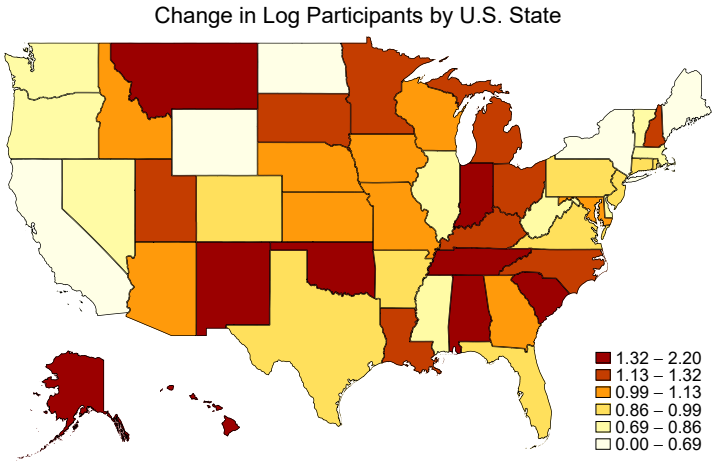

Note: This figure plots the change in the log of the number of robo participants from each U.S. state. The change is from the pre-reduction period (December 1, 2014 to July 7, 2015) to the post-reduction period (July 7, 2015 to February 29, 2016).

Figure A2: Parallel Trends in Robo Participation. Raw Data.

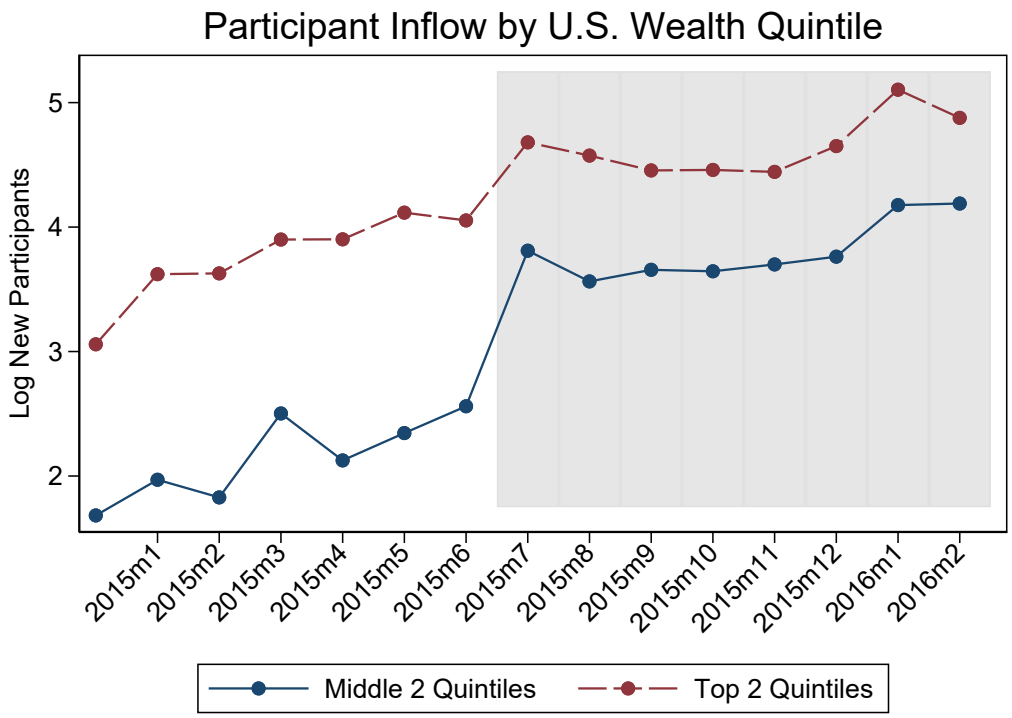

Note: This figure plots the log number of new robo participants from the second and third U.S. wealth quintiles and the fourth and fifth U.S. wealth quintiles, by month. The remaining notes are the same as in Figure 2.

Figure A3: Participants across Robo Advisors. Assessing Reallocation.

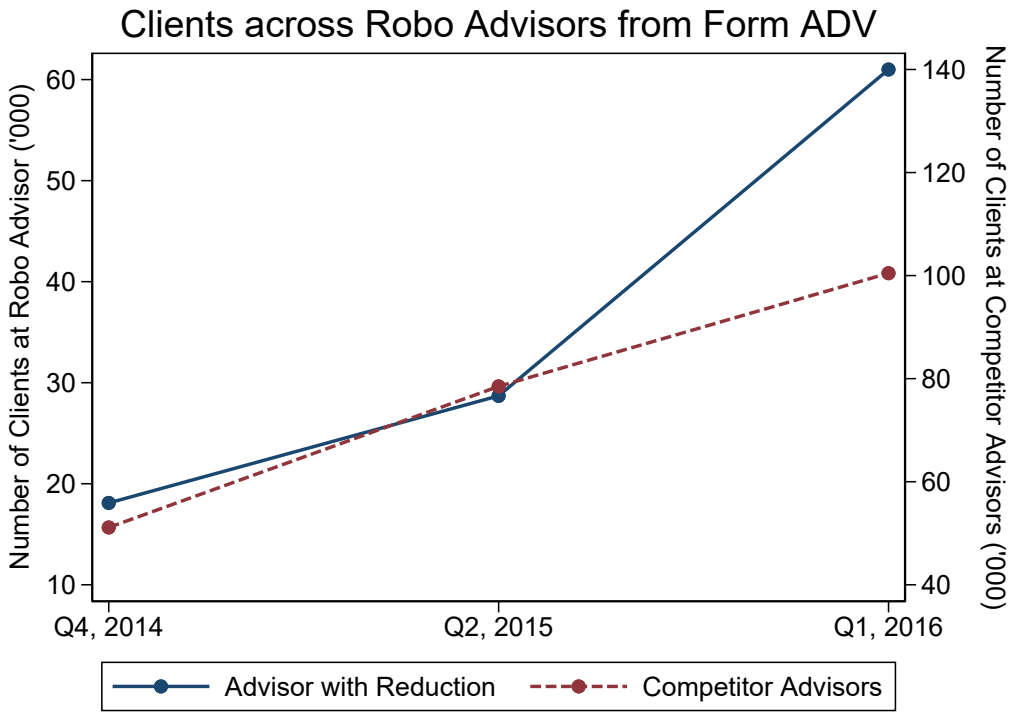

Note: This figure plots the number of clients across robo advisors, in thousands, which assesses whether the reduction increases robo participation or simply reallocates robo participants across advisors. We plot this value separately for the robo advisor that reduced its account minimum, Wealthfront, and for its competitors combined. Data are from the SEC's Form ADV. Competitors are defined as Betterment and Personal Capital, since Schwab's and Vanguard's robo advising services do not file a separate Form ADV. The SEC defines clients to include investors who have not compensated their advisor. Advisors do not file a Form ADV every quarter, and so we use the nearest available observation when the advisor does not file a form ADV in a quarter. We take the average number of clients in cases where a robo advisor files more than one ADV in a given quarter.

Figure A4: Robustness of the Change in Representativeness

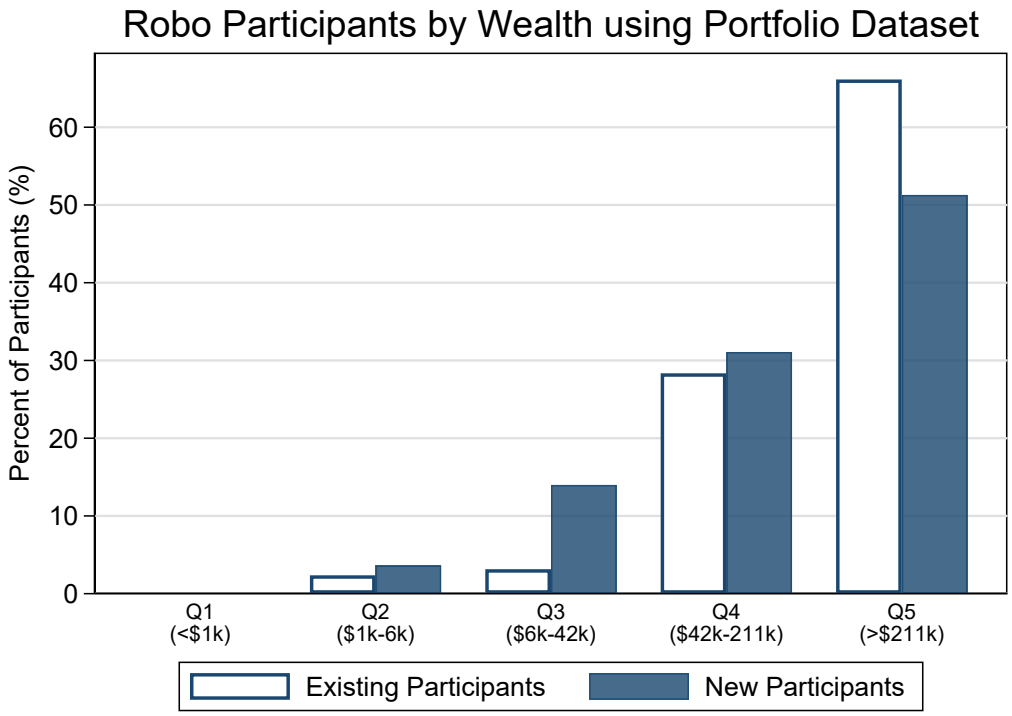

Note: This figure replicates Figure 2 using the subset of robo participants who are also in the portfolio dataset. The remaining notes are the same as in Figure 2.

Figure A5: Share of U.S. Households in the Middle Class by Age

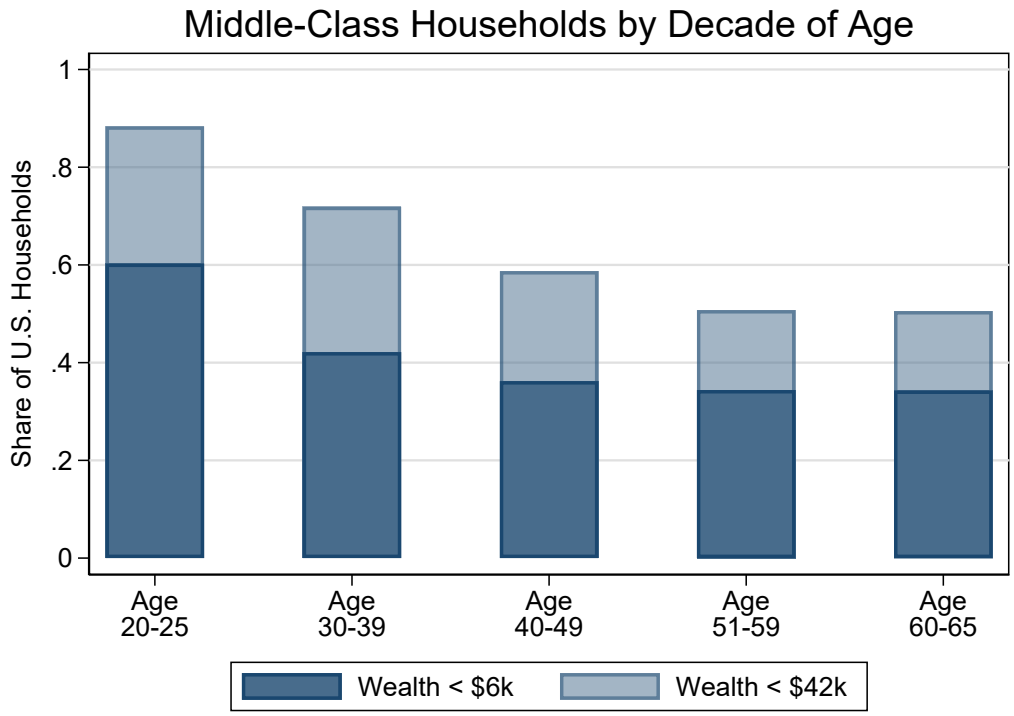

Note: This figure plots the share of U.S. households in the first or second quintile of the overall U.S. distribution of liquid assets (under $6k) and also the share in the first, second, or third quintiles (under $42k). The shares are calculated for each decade of household age. Data are from the SCF dataset.

Figure A6: Model Policy Rules under Previous Minimum

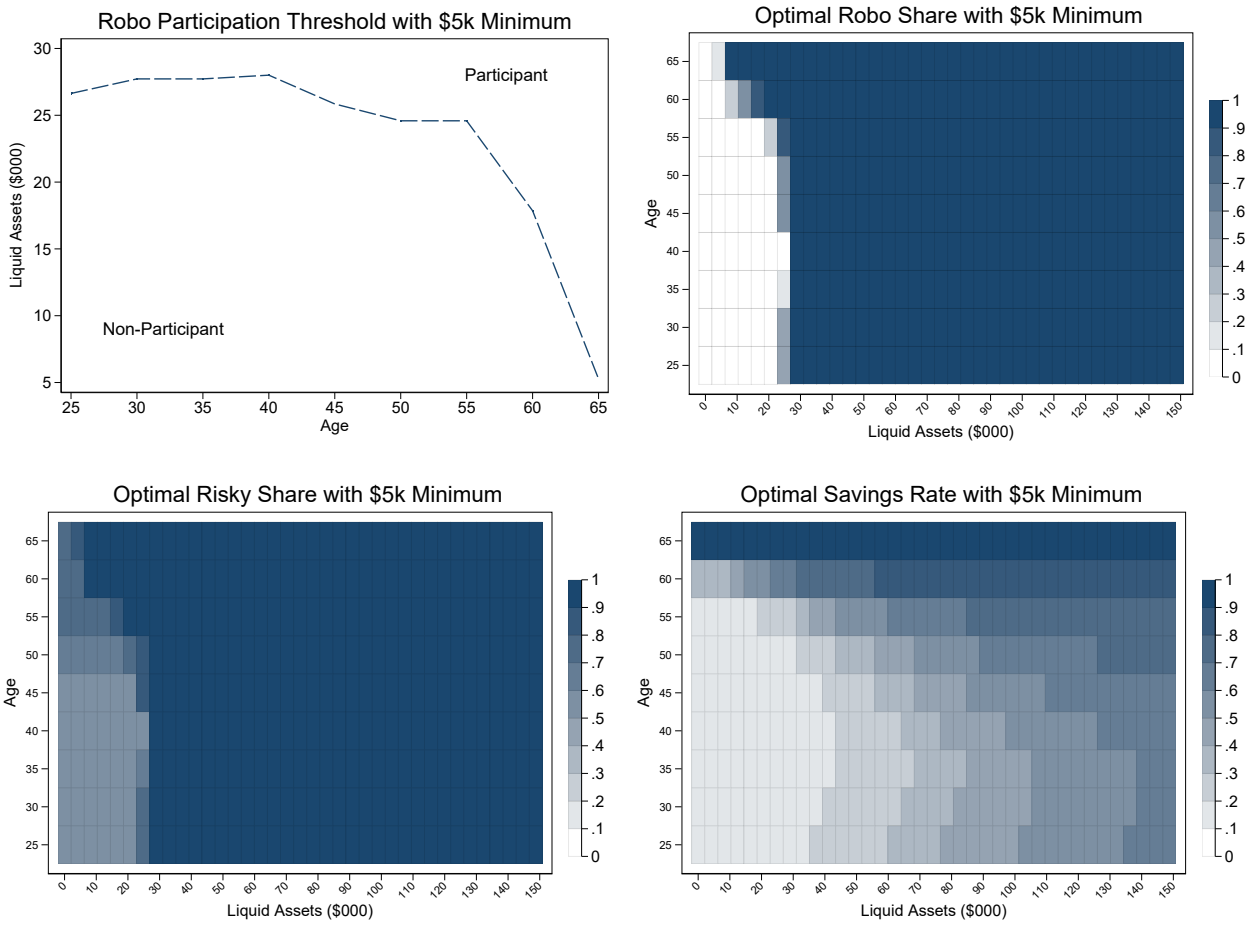

Note: This figure summarizes households' optimal participation, portfolio choice, and savings behavior (i.e., policy rules) under a $5,000 account minimum, based on the model in Section 6. The upper-left panel plots the threshold of liquid assets at which the household participates with the robo advisor, by 5-year age bin. The upper-right panel plots the share of the household's investible wealth allocated to the robo portfolio, where investible wealth equals the difference between cash-on-hand and consumption. The lower-left panel similarly plots the share of the household's investible wealth allocated to either the robo portfolio or the self-managed portfolio. The lower-right panel plots the share of the household's cash-on-hand that is saved. Policy rules are averaged across years in each age bin. Appendix C contains discussion and details.

Figure A7: Robo Value Exposure by Age and Wealth

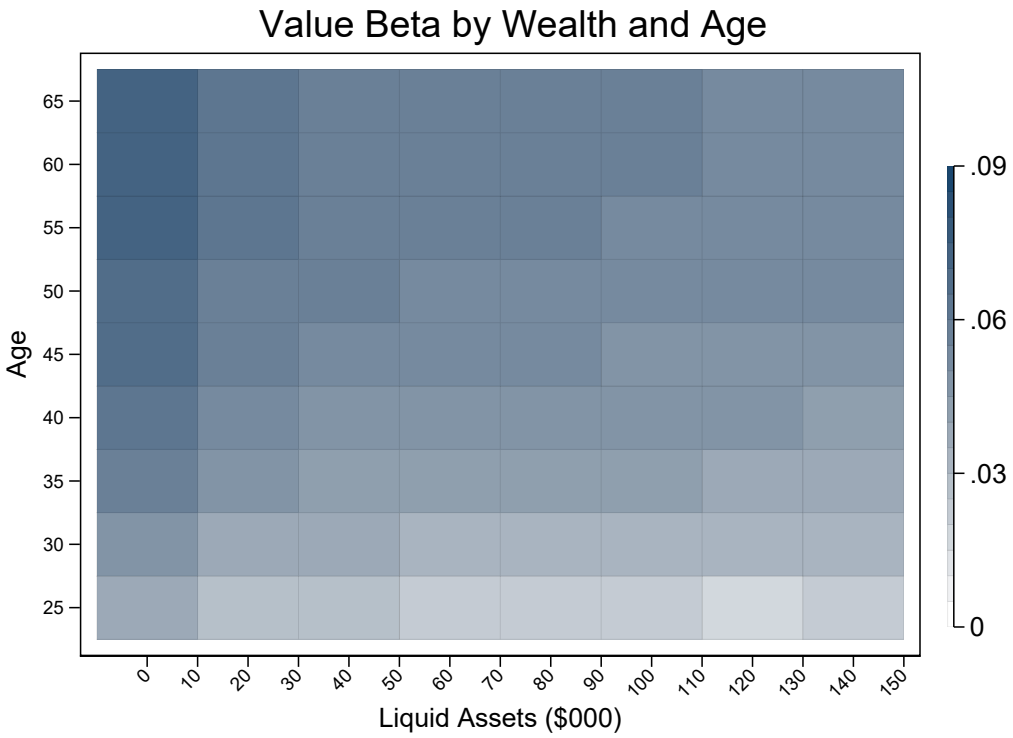

Note: This figure is analogous to Figure 4, and it shows the loading on the High-minus-Low factor (Value Beta). The remaining notes are the same as in Figure 4.

Table A1: Summary of the U.S. Robo Advising Market around the Reduction

Robo advisor	AUM ($bil)	Fees by Account Size	Account Minimum	Presence of Human Advisor
Wealthfront	2.43	0% (under $10k) 0.25% (over $10k)	$500	No
Betterment	2.33	0.35% or $36 (under $10k) 0.25% ($10k to $100k) 0.15 % (over $100k)	$0	Yes (2017)
Personal Capital	1.44	0.89% (under $1mil) 0.49% to 0.89% (over $1mil)	$100k	Yes (2009)
Charles Schwab, Intelligent Portfolios	3	0% (see note)	$5k	Yes (2017)
Vanguard, Personal Advisor Services	21.2	0.3%	$50k	Yes (2015)

Note: This table presents information about the five largest robo advisors in the U.S. market around the time of Wealthfront's reduction in account minimum in July 2015. AUM denotes assets under management around July 2015, which we obtain from the Q2, 2015 Form 13-F for Wealthfront, Betterment, and Personal Capital and from company press releases for Schwab and Vanguard. Fees denotes annual management fees in July 2015, which we obtain from company press releases and contemporaneous industry publications. Fees do not include expense ratios on ETFs in the robo portfolio. Betterment charged 0.35% on accounts under $10,000 which auto-invest at least $100 per month, or $3 monthly (i.e., $36 annually) if they do not auto-invest. Schwab's robo advising service does not charge a management fee, and it instead monetizes by holding 8-10% of clients' portfolios in cash. Account Minimum denotes the account minimum required to open an account in July 2015, which we obtain from company press releases and contemporaneous industry publications. Presence of Human Advisor denotes whether the advisor offers the option to speak with a human advisor, which we obtain from company websites, industry publications, and phone calls with company representatives. The year when the option to speak with a human became available is listed in parentheses. Wealthfront, Betterment, Personal Capital, Schwab, and Vanguard respectively held $23, $22, $13, $45.9, and $179.7 billion in June 2020. Collectively, these five advisors held $283.6 billion in AUM in June 2020, compared to $30.4 billion in July 2015.

Table A2: Summary of Robo Portfolios

Risk Tolerance (0.5 to 10) (1)	CAPM Beta (2)	Stocks (%) (3)	Bonds (%) (4)	Other (%) (5)	Percent of Households (%) (6)	Average Age (7)
0.50	0.32	33.00	60.00	7.00	0.67	39
2.00	0.45	47.00	48.00	5.00	0.39	46
2.50	0.49	50.00	44.00	6.00	0.20	48
3.00	0.52	53.00	41.00	6.00	0.89	48
3.50	0.57	59.00	35.00	6.00	0.86	46
4.00	0.58	59.00	35.00	6.00	1.56	39
4.50	0.61	62.00	33.00	5.00	1.14	42
5.00	0.64	66.00	29.00	5.00	1.68	42
5.50	0.67	69.00	26.00	5.00	1.21	48
6.00	0.70	72.00	23.00	5.00	2.27	40
6.50	0.72	74.00	21.00	5.00	2.32	42
7.00	0.75	77.00	18.00	5.00	6.41	36
7.50	0.77	80.00	15.00	5.00	8.06	39
8.00	0.79	82.00	13.00	5.00	14.39	33
8.50	0.82	86.00	9.00	5.00	16.50	34
9.00	0.85	89.00	6.00	5.00	16.30	33
9.50	0.88	90.00	5.00	5.00	5.43	35
10.00	0.91	90.00	5.00	5.00	19.72	31

Note: This table summarizes robo portfolios assigned to households in our sample. Portfolios are indexed by risk tolerance score, which ranges from 0.5 to 10 in increments of 0.5, and tax status. Each portfolio has a unique vector of weights across 10 possible ETFs, which are chosen to represent exposure to different asset classes. Stocks, Bonds, and Other respectively denote the sum of weights for ETFs that track stock market indices (VIG, VTI, VEA, VW), bond market indices (LQD, EMB, MUB, SCHP), and other asset classes, namely real estate (VNQ) and commodities (XLE). Beta is based on the CAPM, as described in Appendix D. Column (6) shows the percent of robo participants with the indicated portfolio. Column (7) shows the average age of participants with the indicated portfolio. The table only shows taxable portfolios to emphasize how the allocation varies across risk scores, rather than tax status.

Table A3: Summary of U.S. Wealth Quintiles

Wealth Quintile:	First	Second	Third	Fourth	Fifth
Participation in the Stock Market (%)	0.3%	6.4%	31.4%	57.9%	87.0%
Participation with Professional Assistance (%)	0.2%	4.1%	20.7%	41.8%	69.3%
Range of Liquid Assets ($000)	[0,0.6]	[0.6,6.3]	[6.3,42]	[42,211]	>211

Note: This table summarizes the share of U.S. households who participate in the stock market and in asset management by wealth quintile in 2016, based on the SCF dataset. The first row summarizes participation in the stock market, defined as owning stocks, mutual funds, a trust, or an IRA. The second row summarizes participation in the stock market with professional assistance, defined as both participating in the stock market and consulting with a broker, financial planner, banker, accountant or lawyer regarding investment. The bottom row summarizes the range of liquid assets that define each U.S. wealth quintile, in thousands of dollars. Wealth consists of liquid assets, defined as the sum of checking accounts, savings accounts, certificates of deposit, cash, stocks, bonds, savings bonds, mutual funds, annuities, trusts, IRAs, and employer-provided retirement plans. The sample consists of all households in the 2016 SCF.

Table A4: Robustness of Democratization to State-Clustered Standard Errors

$Y_i =$	*New Participant*$_i$					*Under Minimum*$_i$		*At Minimum*$_i$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Middle$_i$	0.219	0.151	0.137	0.145	0.155	0.294		0.253	
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)		(0.000)	
Second Quintile$_i$							0.555		0.309
							(0.000)		(0.000)
Third Quintile$_i$							0.269		0.248
							(0.000)		(0.000)
Middle$_i \times$ *New Participant*$_i$								-0.316	
								(0.000)	
Second Quintile$_i \times$ *New Participant*$_i$									-0.467
									(0.000)
Third Quintile$_i \times$ *New Participant*$_i$									-0.302
									(0.000)
New Participant$_i$								-0.149	-0.149
								(0.000)	(0.000)
Measure of Middle	Second or Third Quintile			Second Quintile	Middle with Buffer	Second or Third Quintile			
Controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.033	0.067	0.097	0.078	0.098	0.156	0.165	0.096	0.097
Number of Observations	9,349	9,349	9,349	7,530	8,982	5,088	5,088	6,890	6,890

Note: P-values are in parentheses. This table reestimates the main specifications in Tables 2 and 3 after clustering standard errors by state of residence.

Table A5: Magnitude of Effect on Robo Participation

	Growth in Number of Robo Participants					
	All Participants			Middle-Class Participants		
	Data (g)	No Reduction (g^C)	Effect (η)	Data (g)	No Reduction (g^C)	Effect (η)
	(1)	(2)	(3)	(4)	(5)	(6)
Baseline Estimates:						
Table 2, Column (3)	119.4%	106.0%	13.4%	239.4%	131.6%	107.8%
Additional Estimates:						
Table 2, Column (6)	119.4%	117.6%	1.8%	301.0%	153.5%	147.5%
Table 2, Column (7)	119.4%	105.6%	13.8%	256.2%	129.7%	126.5%

Note: This table summarizes the observed and counterfactual growth rates in the number of robo participants around the reduction, which assesses the magnitude of the results in Table 2. Column (1) summarizes the observed growth rate in the total number of robo participants, denoted g, and column (2) summarizes the counterfactual growth rate in the absence of the reduction, denoted g^C and defined in equation (B4). Column (3) summarizes the effect of the reduction, defined as the difference between g and g^C, that is, $\eta = g - g^C$. Columns (4)-(5) summarize the analogous observed and counterfactual growth rates in the number of middle-class robo participants, and column (6) summarizes the analogous value of η. Each row calculates these statistics using the estimated coefficient μ and definition of $Middle_i$ from the indicated specification in Table 2, using the methodology described in Appendix Section B.2. The observed growth rate in the number of middle-class participants differs across specifications in column (4) because the definition of $Middle_i$ varies across specifications. The remaining notes are the same as in Table 2.

Table A6: Summary of Non-Robo and Robo Portfolio Loadings

	Middle Class ($N = 354$)			Upper Class ($N = 1,559$)		
	Self-Managed	Matched Robo	Difference	Self-Managed	Non-Robo	Difference
	(1)	(2)	(3)	(4)	(5)	(6)
Factor Loadings (β_i)						
Market	0.930	0.893	-0.036	0.899	0.876	-0.023
			(0.002)			(0.000)
SMB	0.044	0.003	-0.040	0.030	-0.001	-0.032
			(0.004)			(0.000)
HML	-0.086	0.061	0.147	-0.074	0.061	0.135
			(0.000)			(0.000)
GLB	0.629	-0.020	-0.649	0.574	-0.023	-0.597
			(0.000)			(0.000)
USB	-0.447	0.508	0.955	-0.391	0.512	0.903
			(0.000)			(0.000)

Note: P-values are in parentheses. This table summarizes the factor loadings for self-managed and matched robo portfolios, based on the Fama-French Three Factor Model augmented with U.S. and global bond returns. Subscript i indexes portfolio. Each row summarizes the loading on a different factor: Market is the return on the CRSP Value-Weighted Index, net of the risk-free rate; HML is the spread in monthly return between high book-to-market stocks and low book-to-market stocks; SMB is the spread in monthly return between stocks with a small market capitalization and a big market capitalization; GLB is the monthly return on the Barclays Global Aggregate Bond Index Unhedged, net of the risk-free rate; and USB is the monthly return on the Barclays Aggregate U.S. Bond Index Unhedged, net of the risk-free rate. The remaining notes are the same as in Table 6.

Table A7: Covariances and Means of Pricing Factors

Panel (a): Covariance Matrix

	Market	SMB	HML	GLB	USB
Market	0.022	0.005	-0.004	0.001	0.000
SMB	0.005	0.011	-0.002	-0.001	-0.001
HML	-0.004	-0.002	0.009	0.001	0.001
GLB	0.001	-0.001	0.001	0.003	0.001
USB	0.000	-0.001	0.001	0.001	0.001

Panel (b): Means

Market	SMB	HML	GLB	USB
0.076	0.021	0.042	0.060	0.062

Note: This table summarizes the covariance matrix and mean of the baseline asset pricing factor vector, defined in Appendix Table A6. Panel (a) summarizes the covariance matrix. Panel (b) summarizes the mean. Each value in the table is calculated based on the longest available time series over 1960-2017. Over the 2010-2017 period, the volatility of the market factor equals 12.3% and the mean equals 12.6%. The remaining notes are the same as in Appendix Table A6.

Table A8: Summary of Self-Managed and Robo Portfolios by Factor Model

	Middle Class ($N = 354$)			Upper Class ($N = 1,559$)		
	Self-Managed	Matched Robo	Difference	Self-Managed	Matched Robo	Difference
	(1)	(2)	(3)	(4)	(5)	(6)
(a) Sharpe Ratio						
Fama-French	0.366	0.516	0.150	0.370	0.517	0.146
			(0.000)			(0.000)
CAPM	0.334	0.442	0.108	0.339	0.439	0.100
			(0.000)			(0.000)
(b) Expected Return						
Fama-French	0.064	0.071	0.007	0.062	0.069	0.007
			(0.000)			(0.000)
CAPM	0.058	0.063	0.006	0.056	0.062	0.006
			(0.000)			(0.000)
(c) Total Volatility						
Fama-French	0.213	0.137	-0.077	0.197	0.134	-0.063
			(0.000)			(0.000)
CAPM	0.212	0.143	-0.068	0.198	0.142	-0.057
			(0.000)			(0.000)
(d) Idiosyncratic Volatility						
Fama-French	0.151	0.036	-0.115	0.142	0.035	-0.106
			(0.000)			(0.000)
CAPM	0.173	0.079	-0.094	0.160	0.079	-0.081
			(0.000)			(0.000)
(e) Return Loss from Difference in Sharpe Ratio						
Fama-French			0.046			0.039
CAPM			0.035			0.030

Note: P-values are in parentheses. This table assesses the robustness of Table 6 by summarizing self-managed and counterfactual robo portfolios under different asset pricing factor models: the CAPM; and the Fama-French Three Factor Model (Fama-French). The remaining notes are the same as in Table 6.

Table A9: GMM Estimates of Preference Parameters

Parameter	Estimate	
	(1)	(2)
Coefficient of Relative Risk Aversion (γ)	5.1	7
	[TBD]	[TBD]
Discount Factor (δ)	0.92	0.97
	[TBD]	[TBD]
Stock Market Participation Cost (κ)	168.75	206.25
	[TBD]	[TBD]
Weights	GMM Optimal	Equal

Note: This table estimates the model's preference parameters using the generalized method of moments estimator described in Appendix C. The 9×1 moment vector consists of: a 4×1 vector containing the share of robo participants from the second through the fifth U.S. wealth quintiles under the previous account minimum; an analogous 4×1 vector containing the shares under the reduced minimum; and the coefficient μ from an uncontrolled specification of the regression in column (1) of Table 2. Column (1) estimates the parameters when weighting each empirical moment by its inverse standard error (GMM Optimal). For this reason, the share of participants from the first U.S. wealth quintile is excluded from the moment vector, as this share has no empirical variance. Column (2) estimates the parameters when equally weighting each empirical moment. Standard errors are bootstrapped. Brackets correspond to 95% confidence intervals under an asymptotically normal distribution.

Table A10: Sample Selection from Consulting an Online Financial Advisor

$Y_i =$	*Online Advice*$_i$	*New Participant*$_i$	
	(1)	(2)	(3)
Middle$_i$	0.011	0.153	0.125
	(0.000)	(0.000)	(0.000)
log(*Income*$_i$)	0.239	0.004	0.047
	(0.000)	(0.000)	(0.000)
Age$_i$	-0.021	-0.127	-0.613
	(0.000)	(0.000)	(0.000)
Inverse Mills Ratio$_i$			-3.323
			(0.000)
Constant	-1.773	1.840	8.105
	(0.000)	(0.000)	(0.000)
Estimator	Probit	OLS	
R-squared		0.065	0.070
Number of Observations	31,240	9,349	9,349

Note: P-values are in parentheses. This table estimates equation (2) after controlling for selection into consulting an online financial advisor. The selection model and the procedure for estimating it are similar to Heckman (1979). First, using the SCF dataset, we estimate a probit regression that predicts whether a household seeks financial advice online (*Online Advice*$_i$). This variable is the best proxy for whether a household consults a robo advisor among variables in the SCF dataset. Column (1) estimates this probit regression. The variables used in prediction are those observed in both the SCF dataset and in the robo advisory dataset. Observations in column (1) are weighted by SCF sampling weights. Using the coefficient estimates from column (1), we calculate the inverse Mills ratio for each household in our robo advisory dataset (*Inverse Mills Ratio*$_i$). The inverse Mills ratio is proportional to the expected value of unobserved determinants of the probability of robo participation for households who seek financial advice online (Heckman (1979)). Column (2) establishes a benchmark by reestimating the baseline specification from Table 2 after only including control variables observed in the SCF dataset. Column (3) controls for *Inverse Mills Ratio*$_i$, following Heckman (1979). The sample in column (1) consists of the SCF. The sample in columns (2)-(3) is as in Table 2. Column (1) estimates the first stage: *Online Advice*$_i$ indicates whether the household seeks financial advice on the internet. Observations in column (1) are weighted using SCF sampling weights. Standard errors in column (3) are bootstrapped, and the p-value is calculated under an asymptotically normal distribution. The remaining notes are the same as in Table 2.