Title
Robo advisors and access to wealth management

Authors
Reher, Michael
Sokolinski, Stanislav

Publication Date
2024-05-01

DOI
10.1016/j.jfineco.2024.103829

Robo Advisors and Access to Wealth Management*

Michael Reher† and Stanislav Sokolinski‡

January 2024

Abstract

We investigate how access to robo-advisors impacts the financial investment and welfare of less-wealthy investors. We leverage a quasi-experiment where a major U.S. robo-advisor significantly expands access by reducing its account minimum, increasing participation by middle-class investors but not the poor. A benchmark model calibrated to portfolio-level data rationalizes this increase: middle-class investors want sophisticated investing but cannot achieve it themselves. Their welfare rises moderately, driven by advanced features like multi-dimensional glide-paths and additional priced risk factors. Middle-age investors gain three times more than millennials. Our results reveal novel margins of demand for robo-advisors, helping explain their sustained growth.

Keywords: FinTech, Financial Advice, Portfolio Delegation, Inequality

JEL Classification: G11, G24, D3, O3

*We thank Christoph Aymanns, John Campbell, Michael Boutros, 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, Nikolai Roussanov, Andrei Shleifer, Celine Sun, Allan Timmermann, Ross Valkanov, Boris Vallee, Ken Wilbur, anonymous referees, and seminar participants at UC Irvine Finance Conference, Michigan State University, University of Houston, Indiana University, University of Texas at Dallas, CEAR-RSI Household Finance Workshop, 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

‡Broad College of Business, Michigan State University. Email: sokolins@msu.edu

"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 Magazine

1 Introduction

Recent technological innovation has enabled financial intermediaries to scale the provision of many traditional services. In the wealth management market, this scaling has taken the form 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 rapidly in both popularity and size over the past decade. For example, the largest five robo advisors have grown roughly tenfold, and all of the Big-4 U.S. banks now offer a robo advising service.[1]

Robo advisors claim to have two key advantages over their traditional counterparts. First, by adhering to a transparent investing algorithm, robo advisors can benefit their clients by circumventing inefficiencies documented among financial advisors, many of which stem from portfolio-by-portfolio discretion.[2] Second, because automation lowers per-portfolio management costs, robo advisors can extend the benefits of professional wealth management to less-wealthy investors, who do not have enough assets to invest with traditional managers. Our paper examines the latter claim, which is popular and intuitive but has not been studied. We use a quasi-experiment and a quantitative model to trace out the effect of access to robo advisors down to the welfare of new, less-wealthy robo investors.

In the first part of the paper, we study how access to robo advisors affects participation by the less-wealthy. Not only is this exercise important in its own right, but the results will also help discipline the model that we develop in the second part of the paper. However, identifying the effect of interest is challenging for two reasons: lack of standardized data on robo participants; and a need for a setting in which robo advisors suddenly become more accessible to the less-wealthy. We overcome these challenges by obtaining a novel dataset directly from a major U.S. robo advisor and by studying a quasi-experiment. In our setting, described in Sections 2 and 3, the same robo

[1] The top five robo advisors managed \$283 billion in 2020 versus \$30.4 billion in 2015 (Appendix Table A1).

[2] A partial list of inefficiencies includes 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.

advisor suddenly reduces its account minimum from $5,000 to $500. This shock represents one of the first examples in which sophisticated wealth management becomes available to a wide range of less-wealthy investors.

In Section 4, we find that the reduction democratizes the market for robo advisors by bringing in new, middle-class investors. The wealth distribution of robo participants shifts sharply leftward after the reduction, while showing no pre-trend in the prior months. The share of participants from the second and third U.S. wealth quintiles (the "middle class") increases by 107% (16 pps). This increase reflects a sharp break from trend in the number of middle-class participants, whereas there is no such change in the number of participants from the upper two quintiles (the "upper class"). However, the democratization is incomplete, in that we find no change in participation among the bottom quintile (the "lower class"), as conjectured by Philippon (2019).

According to our theory, a relaxation of minimum-account constraints drives this shift in the robo wealth distribution. We sharpen this interpretation through a difference-in-difference (DiD) analysis. The middle class represents the treatment group, since it experiences a relaxation of minimum-account constraints due to the reduction. The upper class is the control group. As a benchmark, we find that middle-class investors are 14 pps more likely to participate with the robo advisor after the reduction, relative to the upper class. We obtain similar results from dynamic DiD designs at various time frequencies, and we perform an event study that strongly supports the assumption of parallel trends.

Section 5 further tests that the empirical effect works by relaxing minimum-account constraints. We first show that the majority of new, middle-class robo participants bunched their investment at the previous minimum of $5,000 prior to the reduction. Such bunching is a hallmark of binding constraints. After the reduction, the bunching immediately disappears, and most new middle-class participants make a previously infeasible investment of under $5,000. Additionally, we find no evidence that the results are driven by other channels such as: heterogeneous response to targeted or non-targeted advertising; gambling motives; business stealing from competitors; heterogeneous trends by demographics or risk attitude; or measurement error in self-reported wealth.

In the second part of the paper, we develop and calibrate a quantitative model of asset allocation with endogenous portfolio delegation and an account minimum. This exercise addresses two major questions that cannot be answered by the reduced-form DiD analysis from the first part. First, we ask whether differences in asset allocation suffice to explain the increase in middle-class participation with the advisor, or whether we must appeal to non-financial channels, like peace-of-mind from portfolio delegation (e.g., Gennaioli, Shleifer and Vishny (2015)). Second, we assess the welfare gain from access to robo advisors, using the standard lifetime consumption metric (e.g., Gomes (2020)). This approach lets us examine the potential benefits of robo advisors more comprehensively than, say, using realized returns, as it accounts for the investor's horizon and human capital.[3] We are particularly interested in how these potential benefits may vary across investors and which features of robo portfolios drive the results.

Our comparison group consists of retail investors who would otherwise manage their risky assets on their own without access to a robo advisor ("self-managed portfolios"). Section 6 uses portfolio-level microdata to document three advantages that robo portfolios have over self-managed ones. 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. Second, they are much better diversified. Third, robo portfolios exhibit personalization by both age and wealth, which we call a "double glide path": exposure to stocks falls as an investor grows older, holding wealth fixed; but it rises as an investor becomes wealthier, holding age fixed. So, robo portfolios are more personalized than target date funds (TDF), which offer a single glide path by age.

In Section 7, we embed these portfolio characteristics in the model and reproduce the empirical results from the first part of the paper. The model closely matches the effects of the reduction in minimum on the wealth distribution of robo participants. Thus, simply accounting for how robo advisors invest differently than retail investors on their own explains the reduced-form evidence quite well, without a clear need for non-financial channels. The economic mechanism that generates these results within the model is precisely a relaxation of minimum-account constraints, with the following key trade-off. Absent the minimum, all investors would like to invest with the robo advisor because of the aforementioned advantages. However, the account minimum requires a larger-than-optimal risky share and savings rate for less-wealthy investors. As a result, these investors prefer to manage their own portfolios, even though they cannot do so as well as the robo advisor. Reducing the minimum relaxes this constraint, prompting investors with low-to-intermediate wealth (middle class) to invest with the advisor. However, lower-class investors still find the required risky investment too large, and so they choose not to participate.

We next examine the magnitude, drivers, and distribution of welfare gains from access to the

[3] Recent and contemporaneous studies have found that robo portfolios have characteristics that would improve welfare in a static model of asset allocation, such as better diversification (e.g., D'Acunto, Prabhala and Rossi (2019), Rossi and Utkus (2021b), Loos et al. (2020)). Our approach builds on these insights by both quantifying the gains from these static features in terms of lifetime utility and by highlighting dynamic benefits (e.g., age glide paths).

robo advisor. The reduction raises welfare of new robo participants by 0.8%, in terms of lifetime consumption. This increase is meaningful yet plausible.[4] Interestingly, we find substantial heterogeneity by age: the reduction raises welfare three times as much for investors over age 55 (1.7%) than for those under age 35 (0.6%). This difference arises from differences in lifetime human capital, mainly because younger investors have longer working lives. As a result, many of them would have accumulated enough earnings to overcome the previous minimum anyway, whereas older investors would not have become robo participants without the reduction.

Since our comparison group consists of self-managed portfolios, the welfare gains for new robo participants must fundamentally reflect the different characteristics of robo portfolios relative to self-managed ones. We evaluate which characteristics add the most value by calculating the welfare gain for counterfactual robo portfolios after removing each of their advantages, one at a time. Without improved diversification, the welfare gain is 70% lower. Similarly, removing exposure to bond risk factors reduces the welfare gain by 70%. By contrast, differences in exposure to the overall stock market have almost no impact on welfare. This finding reflects how investors in our data attain such exposure on their own. Interestingly, we find that personalization by wealth improves welfare as much as does personalization by age, highlighting the importance of a "double glide path".

Taken together, our results present a novel perspective on the impact of robo advisors. First, robo advisors do not solely impact affluent investors, who already have access to professional managers: robo advisors also expand access to the less-wealthy by relaxing minimum-account constraints. Removing constraints raises participation because less-wealthy investors have strong demand for sophisticated wealth management. This result challenges the canonical model of a self-sufficient investor. Our results also challenge models in which non-financial benefits drive the demand for delegated management, since we can explain this demand as a rational response to improved asset allocation. In that vein, our evidence matches survey results from Rossi and Utkus (2021a), who find that investors seek robo advisors more to improve their investment ability and performance than to achieve peace of mind. More broadly, we exemplify how sophisticated wealth management is not necessarily a luxury good.

The second contribution of this paper is to rigorously document the properties of a fully-

[4] Measuring welfare gains as a percent of lifetime consumption has a long tradition in the literature dating back to at least Lucas (1987). A gain of greater than 0.5% is typically considered consequential. For reference, a gain of 0.8% lies well-within the range of gains from correcting various investment mistakes in a workhorse model developed by Cocco, Gomes and Maenhout (2005).

automated asset allocation rule and to assess which of its features investors value most. We show that robo advisors do not simply substitute for a low-cost equity index fund or a TDF. Instead, much of their value comes from providing sophisticated features like a double glide path and exposure to multiple risk factors. Less-wealthy investors appreciate these features because they struggle to enact similar features on their own. Older investors especially benefit because they rely less on non-financial income, a finding that challenges robo advisors' stereotype as a product for millennials. Overall, our results help explain the sustained growth of the robo market and its integration with the traditional financial sector, despite the many other retail products available during the FinTech era.

Our conclusions are based on the subset of retail investors who are interested in fully automated wealth management, rather than the average U.S. investor. The investors in our sample are presumably more sophisticated than the general U.S. population, since, for example, many of them at least participated in the stock market before joining the advisor. However, a unique feature of our setting is that the advisor's algorithm allocates the same robo portfolio to all investors with the same demographic profile and risk attitude, regardless of sophistication. By adhering to this form of algorithmic fairness, the robo advisor plausibly generates higher gains for less-sophisticated investors who either do not participate in the stock market or do so with major pitfalls. Therefore, our welfare calculations would constitute a lower bound with respect to the gains for the general population. By the same logic, our estimated effect on participation would constitute an upper bound.[5]

Related Literature

This article speaks to a nascent literature on FinTech intermediaries, which includes robo advisors. Unlike prior and contemporaneous studies, we study 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 (2022); D'Hondt et al. (2020)); no option for human interaction (e.g., Rossi and Utkus (2021b)); and robo advisors unaffiliated with the banking system (e.g., Loos et al. (2020)). We also focus on less-wealthy investors in a quasi-experimental setting, complementing the literature's typical focus on more-affluent investors.[6]

[5]Recent evidence suggests that retail investors prone to particular behavioral biases endogenously use FinTech products that perpetuate those biases (e.g., Cookson, Engelberg and Mullins (2023); Barber et al. (2022); Ben-David et al. (2022)). These results would imply that bias-prone investors prefer not to participate with robo advisors because they dislike the "textbook" aspect of robo portfolios.

[6]Our study also relates to a broader question of how technological innovation affects financial inclusion and wealth

Within the literature on robo advisors, our study shares common ground with Reher and Sun (2019), who examine how deposit inflow changes after the 2015 reduction in minimum. Several crucial distinctions set our work apart. Most importantly, our novel structural approach enables both a rigorous calculation of welfare gains and an assessment of the economic mechanisms driving the surge in account formation after the reduction.[7] In particular, we can challenge and refute three potential mechanisms left open by the atheoretical approach in Reher and Sun (2019). First, we challenge the idea that investors become robo participants purely to follow a popular trend; instead, we quantitatively rationalize their response using differences between robo and self-managed portfolios. Second, we challenge the notion that they respond solely because they cannot form a diversified portfolio of stocks; we show that a hypothetical investor in a low-cost equity index fund would also gain from access to robo portfolios, underscoring the importance of more-sophisticated features like additional risk factors and glide paths. Third, we challenge the conventional assumption that robo advisors only benefit younger investors; while the young comprise a large share of new robo participants, the primary beneficiaries are actually older investors with limited investable wealth. Our empirical approach also differs from Reher and Sun (2019) by rigorously accounting for time-varying confounds, particularly targeted advertising, which enables us to credibly identify the reduction's causal effect. Lastly, Reher and Sun (2019) only document a response by investors with less than \$107,500 in wealth, but we show that the reduction impacts investors much farther down the U.S. wealth distribution, notably those with \$6,000 to \$42,000 (third quintile) and especially those with \$1,000 to \$6,000 (second quintile).

Our study also delivers a new perspective on the drivers of demand for professional portfolio management. The literature has argued that retail investors seek portfolio managers because managers actively collect information about the underlying assets (e.g., Gârleanu and Pedersen (2018)) or confer peace of mind (e.g., Gennaioli, Shleifer and Vishny (2015)). We find that many less-wealthy investors have more mundane needs: to construct personalized and diversified portfolios, even if the underlying assets are passively managed.[8] These mundane needs are not fully met by another form of personalized asset management, TDFs (e.g., Balduzzi and Reuter (2019);

inequality. See, for example, the studies in the contexts of app-based payments (e.g., Hong, Lu and Pan (2022)), bank deposits (e.g., Bachas et al. (2018); Bachas et al. (2020); Higgins (2022)), and mortgage markets (e.g., Fuster et al. (2019); Bartlett et al. (2021); Fuster et al. (2021)).

[7] By studying a model of portfolio choice with a dynamic decision between self-managed and professionally managed portfolios, we also contribute to the literature on quantitative life cycle models summarized by Gomes (2020).

[8] This finding corroborates the Von Gaudecker (2015) result that Dutch investors diversify better when they consult a financial professional. Insofar as wealth correlates with financial literacy, our findings are also consistent with the Bianchi (2018) result that less-literate investors struggle to rebalance to their optimal risk exposure.

Parker, Schoar and Su (2021)), as robo advisors offer more personalization via a double glide path.

Finally, we contribute to the household finance literature by characterizing account minimums as a friction that distorts investment in risky assets. This finding demonstrates the importance of frictions that arise from the supply side, like internet speed (e.g., Hvide et al. (2023)), as distinct from frictions that directly depend on investor characteristics like: 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)).

2 Institutional Background

Below we describe the U.S. robo advising market (2.1), the advisor's service we study (2.2), and our quasi-experiment (2.3). We use the terms "robo advisor" and "automated wealth manager" synonymously throughout the study.[9]

2.1 The U.S. Robo Advising Market

Robo advisors emerged in the mid-2000s in response to the limitations of traditional wealth managers (D'Acunto and Rossi (2020)). They rely on algorithms to select and maintain an allocation for their clients. This 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)). 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 $2.4 billion and was the largest standalone robo advisor in the U.S. market, with Betterment and Personal Capital as its nearest competitors. Wealthfront is also the only robo advisor that relies purely on automation, with no option for a human advisor. Two traditional managers, Vanguard and Charles Schwab,

[9]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"), as opposed to just providing financial advice. This pipeline includes the baseline product that we study (a personalized and automatically rebalanced portfolio) as well as: tax loss harvesting; a risk parity fund; and financial planning on both the asset and liability sides of the investor's balance sheet.

launched robo advising services early in 2015. Both of these services managed more than Wealthfront because they transferred assets from existing, non-robo services. From here on, we refer to Wealthfront simply as "the robo advisor".

2.2 The Robo Advisor's Baseline Service

We study the robo advisor's baseline service: an automatically rebalanced portfolio of 10 ETFs corresponding to 10 asset classes. The 10 ETFs are chosen to track stock market indices (VIG, VTI, VEA, VWO), bond market indices (LQD, PCY, MUB, SCHP), and other asset classes, namely real estate (VNQ) and commodities (XLE). The portfolio weights are determined by a questionnaire about the client's 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.[10]

Importantly, robo portfolios are not recommendations: rather, they are directly selected, managed and rebalanced by the robo advisor. Consequently, investors 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, Campbell and Sodini (2007)), willingness to follow advice (e.g., Bhattacharya et al. (2012)), or reluctance to rebalance (e.g., Calvet, Campbell and Sodini (2009)). For its services, the advisor charges an asset-based management fee of 0% for accounts under $10,000 and 0.25% for larger accounts. So, most of the less-wealthy investors in our data incur no management fee.

2.2.1 Risk Profile of Robo Portfolios

We summarize two important characteristics of robo portfolios that will later explain many of our results. This summary uses the core robo datasets described in Section 3.2.

First, robo portfolios are very well-diversified, with little idiosyncratic risk. To illustrate this

[10] Each asset class has a secondary ETF, which the advisor will occasionally hold instead of the primary ETFs named in the text. Since the robo advisor does not hold fractional shares in its baseline product, it will substitute the secondary ETF when the share price of the primary ETF is sufficiently high, relative to the portfolio's size, that the target portfolio weight cannot be achieved within reasonable tracking error (Wealthfront (2023)). We conduct our main analysis under the approximation that the advisor holds fractional shares, but we will assess robustness to this approximation in Section 6.2.1. Appendix Table A3 lists the secondary ETF for each asset class and reports the correlation in total return with the primary ETF. The correlation coefficient is quite high and often greater than 0.99.

feature, we first separate idiosyncratic risk from priced risk, using a five-factor model described in Section 6.2.1 (three Fama and French (1993) factors and two bond factors). Then, we calculate total volatility across portfolios with different risk tolerance scores as well as the volatility that comes from factor exposure ("volatility from priced risk"). Panel (a) of Figure 1 shows that the difference between the total volatility and the volatility from priced risk is small, implying very little idiosyncratic risk. In addition, this small and constant difference holds across risk tolerance scores, which means that all robo participants receive a well-diversified portfolio.

Second, robo portfolios are personalized along a variety of dimensions, the most important of which are age and wealth. For example, panel (b) shows that they provide less priced risk to both older and less-wealthy investors. These features agree with two recommendations of canonical life cycle models: older investors take less risk because they hold a smaller share of total wealth in human capital; and less-wealthy investors do so to hedge background risk. As a result, robo portfolios offer a "double glide path" by age and wealth, as opposed to TDFs that adjust asset allocation by age only. The robo advisor implements these glide paths by reducing stock exposure and raising bond exposure for older and less-wealthy investors (panels (c) and (d)). Appendix Figure A6 shows that the advisor also raises investors' exposure to value stocks as they age (e.g., Jurek and Viceira (2011)).

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 an investor would need to invest only \$500 to participate with the advisor as opposed to \$5,000 beforehand. This reduction is sizeable from the standpoint of most U.S. retail investors. For reference, \$5,000 equals 30% of median liquid assets (\$17,000), and it defines the 37th percentile of the U.S. wealth distribution, according to the 2016 Survey of Consumer Finances. So, prior to the reduction, half of U.S. retail investors could not participate without investing at least 30% of their wealth, while 37% could not participate at all without borrowing.

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.[11] Importantly, the month of the reduction does not coincide with any other product launches by the robo advisor, any changes in its fee, or

[11] 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.

any significant developments in the overall robo advising market. Since our data are at a relatively high (weekly) frequency, this effectively-random timing allows us to identify the reduction's effect on robo participation.[12]

3 Data

Our analysis relies on two core datasets that draw from the same population: a panel dataset covering deposit activity by investors who participate with the robo advisor (3.1); and a dataset on self-managed, outside portfolio holdings that includes both investors in the deposits dataset and robo non-participants (3.2). Appendix A has details, including information on additional datasets that we use. We use the term "robo participant" to describe investors 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. Straddling the July 2015 reduction, this time window also marks a formative period in the history of the robo advising market when the number of robo participants was still small. For each investor, we observe the size of their weekly deposits and the following demographic variables: annual income; state of residence; 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 are static. We address measurement error from potential misreporting later in Section 4.3. To match the setup of our model, in which account holders have at most one risky account, we treat taxable and non-taxable accounts of the same account holder in our deposits dataset as separate observations (separate "investors"). Section 5.3 reproduces our main results after dropping the 12% of account holders who have multiple accounts.

Studying a private, company-specific dataset has two major 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, Form ADV data includes inactive clients who stopped making deposits as well as "clients" who create a username but never provide the robo advisor with any

[12] We make no claim as to the advisor's motivation or the reduction's optimality as a business decision. In its press releases, the advisor motivates the reduction as a bet that less-wealthy customers will eventually accumulate enough assets to become profitable (Nash (2015)).

capital to manage.[13] Second, our dataset includes information about a robo participant's wealth, allowing us to study investment activity across the wealth distribution.

We use the 2016 Survey of Consumer Finances (SCF) to categorize investors into three wealth groups. We respectively use the terms "lower class", "middle class", and "upper class" to describe investors from the first, second or third, and fourth or fifth quintiles of the U.S. distribution of liquid assets. The resulting boundary between the lower versus middle class is \$1,000, and the boundary between the middle versus upper class is \$42,000.[14] Since the SCF includes a representative sample of U.S. retail investors, this approach ensures that we define wealth groups in accordance with the overall U.S. population.

3.1.1 Summary Statistics

Table 1 summarizes the deposits dataset. Panel (a) compares new participants who join the robo advisor after the reduction in account minimum 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. They are also more likely to come from less financially developed U.S. states (Appendix Figure A1). Importantly, new robo participants invest with persistence, as 98% of them do not close their account over our sample period; and 71% make subsequent deposits that are typically much smaller than their initial investment. These results suggest that new participants did not join the advisor for a short-term experimentation.

Panel (b) conveys similar patterns within the middle-class, with two principal findings. First, 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. Second, new middle-class participants are not significantly younger than existing ones, suggesting the reduction does not disproportionately affect the young. We explore both findings later, in our quantitative model.

3.2 Portfolio Dataset

The second core dataset contains information on self-managed portfolio holdings for a subset

[13] 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).

[14] See Appendix Table A4 for detail. Liquid assets in the SCF are calculated to match the definition in our robo advising dataset as closely as possible: 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.

of investors who use the advisor's online financial advice tool. The tool works by loading data from the investor's brokerage account and providing free financial advice about the portfolio in that account. The tool also shows the investor the portfolio they would receive if they become a client of the advisor. Some investors decide to join the robo advisor, migrating into the deposits dataset. Other investors decide not to become robo participants. The combined dataset comprises 1,913 pairs of self-managed and counterfactual robo portfolios for both robo participants and non-participants. Investors in the portfolio dataset have similar demographic characteristics as those in the deposit dataset (Appendix Table A5).

Three unique features of the portfolio dataset allow us to study an investor's choice between self-management versus delegation to a robo advisor. First, we observe robo non-participants (45% of sample), which avoids selection bias from the possibility that only investors who cannot invest efficiently on their own choose to delegate. Second, we can ensure that outside portfolios are indeed self-managed because we observe advisory and management fees. Accordingly, we restrict our analysis to portfolios without such fees, corresponding to 73% of the overall sample and 85% of the middle class. Third, we directly observe the counterfactual self-managed portfolio for each robo portfolio. Specifying such a counterfactual is critical for measuring the effects of delegated management (Chalmers and Reuter (2020)).

We use both deposit and portfolio datasets throughout our study to accomplish separate goals. We primarily use the deposits data in the first part of our paper, Sections 4 and 5. This dataset contains information on all robo participants, which is important for estimating the effects of accessibility on participation. We primarily use the portfolio dataset to study the drivers of demand for robo advisors and welfare in Sections 6 and 7, since the aforementioned features make it especially suitable for this purpose.

4 Effect of Accessibility on Robo Participation

4.1 Shift in Wealth Distribution of Robo Participants

We first examine how the wealth distribution of new robo participants changes over time. Figure 2 shows that it shifts sharply leftward in the third quarter of 2015, immediately after the reduction. The wealth distribution is stable in all of the preceding quarters, suggesting that the post-reduction shift does not simply reflect a pre-trend. These results imply that reduction "democratizes" the robo market by bringing in new less-wealthy investors.

Figure 3 compares the distributions of new and existing participants across U.S. wealth quintiles. The effect of the reduction is concentrated among the middle-class. 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). This distributional shift reflects an increase in the number of middle-class participants, not a decrease in the number of upper-class participants (Appendix Figure A2). The democratization is incomplete, though, since investors from the first quintile remain non-participants. We find similar results using the subsample of participants in the portfolio dataset (Appendix Figure A3).

We will use the previous graphical evidence to discipline our quantitative model. Accordingly, we must first verify that it indeed represents the effect of the reduction. We use a difference-in-differences (DiD) approach to achieve this purpose.

4.2 Difference-in-Difference Setup

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 investor; $\mathcal{T}$ 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 investor i participates with the robo advisor at some point in period $\mathcal{T}$; $Middle_i$ equals one if i belongs to the second or third U.S. wealth quintile; ζ_i is an investor fixed effect; and X_i is a vector of investor characteristics: age, log income, state of residence fixed effects, and a measure of subjective risk tolerance from the advisor's questionnaire.

Our parameter of interest is μ, which equals the double difference in the probability of becoming a robo participant after versus before the reduction between middle-class versus upper-class investors. We hypothesize that the reduction affects robo participation by relaxing minimum-account constraints. Accordingly, we interpret middle-class investors as the treatment group and upper-class investors as the control group.

The additional terms in equation (1) capture channels distinct from minimum-account constraints. The investor fixed effect ζ_i captures slow-moving characteristics that predispose investors to participating with the advisor, such as sophistication or trust. The interaction $X_i \times Post_{\mathcal{T}}$ captures heterogeneous trends by observed investor characteristics. If, for example, younger investors 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 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 investor i becomes a robo participant after the reduction; and $u_i \equiv \Delta v_i$. The following identification assumption allows us to interpret μ as the causal effect of the reduction:

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

Equation (3) states that unobserved determinants of a *change* in robo participation, u_i, do not systematically vary across the middle and upper classes except because of the reduction in account minimum, conditional on the investor's observable characteristics, X_i. Under this condition, the difference in the change in robo participation between the middle and upper classes reflects the effect of relaxing minimum-account constraints.

There are three ways in which equation (3) could be violated. The first is measurement error in self-reported liquid assets. The second source of bias is an omitted variable that is not caused by the reduction, but it nevertheless affects changes in middle-class investors' robo participation. One such confounding variable could be trend growth in investors' preferences for robo advisors. The third source of bias is an omitted variable that is actually caused by the reduction. The leading example is advertising. If the effects of advertising by Wealthfront are stronger or weaker for middle-class investors, then μ confounds the effect of advertising with the effect of minimum-account constraints. We deal with each of these concerns in our analyses below.

4.3 First-Difference Results

We first estimate equation (2) in the deposits dataset. Therefore, we interpret μ as the reduction's effect, conditional on eventually participating with the robo advisor. One can also interpret μ as the share of robo participants whom the reduction has caused to participate, as Appendix B.3 derives. Table 2 gives the results. Column (1) shows that middle-class investors are 15 pps more likely to become robo participants after the reduction. The robustness of this estimate to state fixed effects (column (2)) implies that the result is not driven by the propensity of middle-class investors to live in certain locations. Neither does the effect vary with age (column (3)), in line with the descriptive evidence in Table 1. We cluster standard errors by investor, but the results remain statistically significant when clustering standard errors by state of residence (Appendix Table A8).

In column (4) we estimate equation (2) using the portfolio dataset. In this case, we need not condition the interpretation of μ on eventual participation because we observe non-participants as well. We find the reduction raises middle-class investors' probability of participation by 10.6 pps in this sample. The slightly smaller effect suggests that middle-class investors who use the online advice tool may have weaker demand for robo advisors.[15]

Our treatment exposure variable, $Middle_i$, may be subject to additive measurement error due to self-reporting. As we show formally in Appendix B.2, such measurement error can bias the estimate upwards only 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 investors from the third quintile from the sample. Under this definition, upper-class investors would need to underreport liquid assets by at least \$36,000 to be misclassified as middle-class. Second, we exclude investors 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 are roughly the same as our baseline estimate.[16]

Lastly, we examine how the effects of the reduction vary within the middle class. We alter our main specification from equation (2) by replacing $Middle_i$ with $Lower\ Middle_i$ (second wealth quintile) and $Upper\ Middle_i$ (third wealth quintile). We estimate a larger effect of the reduction among the lower half of the middle-class, shown in Appendix Table A7.

4.4 Dynamic Difference-in-Difference Results

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

[15] Estimating the effect using the portfolio dataset makes the results dependent on the sample of investors who actively chose to use the online tool. In this case, we interpret μ as the effect conditional on using the tool, rather than on eventually participating. The advantage of using both datasets is that we can test the robustness of our results across different subpopulations of investors.

[16] To sharpen the interpretation of our estimates, we use them to back out the reduction's effect on the overall number of robo participants. Appendix Section B.1 describes this procedure. The results in Appendix Table A10 imply that the reduction increases the number of middle-class participants by 108% and the overall number of robo participants by 13%. Thus, the reduction seems to have spurred significant growth in overall participation with the robo advisor.

higher-frequency dynamic effects influence the baseline results. Accordingly, we estimate the dynamic difference-in-difference (DiD) equation

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

where i and t index investor 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; ζ_i is an investor fixed effect; and ϱ_t is a week fixed effect. Standard errors are clustered by investor and week.

Columns (7)-(8) of Table 2 establish a benchmark effect of between 0.6 pps and 0.8 pps. Compounding this weekly effect over the 32-week post-reduction period implies a cumulative effect that is on par with our first-difference estimates (e.g., 17.6% = $(1-(1-0.006)^{32})$).

The main attraction of a dynamic approach is the ability to evaluate dynamic confounds through an event study. We interact $Middle_i$ in equation (4) with a vector of month fixed effects to trace out the reduction's full dynamic effect,

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

where the coefficients μ_m on the treatment indicator, $Middle_i$, vary non-parametrically by month m, which includes week t. These coefficients represent the difference in outcomes between the middle and upper class in each month. We omit the month prior to the reduction, June 2015, such that we interpret the level of μ_m relative to this base period. The parallel trend assumption requires that $\mu_m = 0$ for all months prior to July 2015.

The results in Figure 4 show that the probability of becoming a robo participant increases sharply and significantly for middle-class investors exactly in the month of the reduction, July 2015. Importantly, the differences between middle and upper-class investors are economically small and statistically indistinguishable from zero before the reduction, providing strong support for the parallel trend assumption. It is thus highly unlikely that pre-trends in middle-class investors' 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 is itself highly unlikely.[17]

[17] Specifically, the absence of a pre-trend rules out bias from the launching of new robo products by the two traditional managers described in Section 2.2. These new products were launched at least two months prior to the reduction (March

5 Robustness and Extensions

The purpose of the following robustness exercises is to distinguish between the channel relevant for our model, namely a relaxation of minimum-account constraints, and alternative explanations, like changes in the advisor's visibility.

5.1 Initial Deposits

If middle-class investors face minimum-account constraints, then they are more likely to make an initial deposit exactly at the prior minimum of $5,000, as previously shown in Table 1. Such "bunching" is a hallmark of constrained behavior. By removing these constraints, the reduction should theoretically reduce bunching at the prior minimum.

We test this hypothesis by replacing the outcome in equation (5) with an indicator that equals one if the initial deposit made by investor i in week t equals $5,000 and zero otherwise.[18] The results in Figure 5 support the role of minimum-account constraints. Consistent with a subsequent relaxation of constraints, middle-class investors who become participants after the reduction are 20-30 pps less likely to invest right at the prior minimum. The differences in the outcome variable between the middle-class and upper-class are statistically indistinguishable from zero before the reduction, in line with the parallel trend assumption.

Appendix Table A9 further shows that this effect is roughly 50% larger for investors from the second wealth quintile than for the third quintile. Since minimum-account constraints should theoretically bind more strongly for the less-wealthy, these differences within the middle-class further support the constraints channel.

5.2 Effects of Advertising

We next examine the potential bias in our results due to changes in the advisor's advertising (ad) policy that coincide with the reduction and differentially affect middle-class invetors. To address this concern, we gather data on the advisor's ad spending during the period around the reduction in account minimum.

and May), comfortably before the strong divergence in middle-class investors' behavior. We reiterate that these two traditional managers largely expanded their robo products among existing clients, and so they effectively compete in a different pool than our advisor.

[18] We include observations where the initial deposit is within 5% of $5,000 to account for potential reporting errors. Excluding these observations does not affect our results. Since the sample consists of investor-weeks in which the investor makes its initial deposit, we cannot include investor fixed effects.

We source our data from Kantar's Vivvix dataset, which contains comprehensive information on advertising spending. Kantar is a prominent business intelligence data provider whose data are commonly used in the literature on the advertising of mutual funds (e.g., Kaniel and Parham (2017); Roussanov, Ruan and Wei (2020)). Details on data construction are outlined in Appendix A. In summary, we compile monthly ad spending data by Wealthfront in 2015 across various geographical areas, media sources, and products. The geographic unit in Kantar is the Designated Market Area (DMA), which we map to counties using the crosswalk from Gentzkow (2006) and Gentzkow and Shapiro (2008). Then, we merge our advertising data with county-level imputed stock market wealth from Chodorow-Reich, Nenov and Simsek (2021). This enables us to analyze ad spending across areas with different wealth levels.

Our analysis focuses on two key questions. Firstly, did the advisor alter its advertising strategy to target investors based on their wealth levels? Secondly, did any such changes impact the likelihood of investors joining the advisor after the reduction? To address the first question, we examine the time-series of ad spending surrounding the reduction. Panel (a) of Figure 6 illustrates monthly ad spending, revealing an ad campaign by the advisor. Notably, the advisor initiated advertising three months before the reduction, reaching a peak monthly spending of \$3 million just before the reduction. Subsequently, the spending decreased to approximately \$1 million post-reduction. The campaign concluded three months after the reduction, with no ad spending in the remaining months of 2015. Importantly, all 2015 ad spending pertained to the advisor's main product analyzed in this study, namely, "Wealthfront Online."

TV advertisements are the dominant method of advertising in the data, accounting for over 99% of spending over the three months before and after the reduction. The small remaining share consists of spending on internet search and display advertisements.[19] Almost all TV ad spending (98%) is on cable TV advertisements. Considering that addressable advertising on cable TV was less common during that period, it seems unlikely that the advisor specifically targeted individual investors. Nonetheless, there remains the possibility that the advisor broadly targeted investors

[19] We check whether the dominance of TV advertisements is due to measurement error in the Kantar dataset by tabulating the same statistic using Nielsen's Ad Intel dataset. The Nielsen dataset has a similar form as Kantar's, but it does not contain information on spending across geographic markets and, so, cannot be used for our main analysis. The share of the robo advisor's advertising spending on TV advertisements over the same time period in the Nielsen dataset also equals 99%, suggesting that the main results based on Kantar are not an artifact. Even if the advisor executed unobserved targeted internet advertisements based on wealth, we can still identify the impact of the reduction under the reasonable assumption that internet ad spending targeted at people with a certain wealth level co-varies with television ad spending targeted at regions where average wealth is the same level, which we do observe. Lastly, the marketing literature does not have a clear consensus about the impact of internet advertising on advertiser revenue, as summarized by Gordon et al. (2020), which suggests that unobserved internet advertisements would have minimal impact on our results.

in specific market areas according to household wealth in those areas.

To investigate this possibility, we categorize a Designated Market Area (DMA) as middle-class if its imputed stock market wealth falls within quintiles two and three of the wealth distribution across all DMAs. DMAs in the top two quintiles are classified as upper-class, mirroring the classification method used for individual investors in our main datasets. In panel (b) of Figure 6, we present the total spending across DMAs, revealing a notable trend: upper-class areas received substantially more ad spending than their middle-class counterparts. This observation suggests that the advisor directed its advertising efforts toward wealthier investors.[20] Appendix Figure A5 reinforces this notion by demonstrating statistically significant differences in average ad spending between middle-class and upper-class areas.

We then delve into examining whether these spending disparities across market areas translated into differences in the likelihood of investors joining the advisor post-reduction. Given that our primary dataset includes information on the investor's home state, we convert DMA-level spending into state-level spending. Our initial measure of state-level spending, denoted as $Adv_{s,t}$, is the average spending across all DMAs in state s in the month of week t. To assess how advertising influenced the response of middle-class investors to the reduction, we extend our DiD equation (4) by including the interactions $Middle_i \times Adv_{s,t}$ and $Middle_i \times Post_t \times Adv_{s,t}$. We incorporate state-by-week fixed effects, which absorb the direct effect of $Adv_{s,t}$ and also account for all state-level unobservables, including differences in advertising costs across states

Our findings, presented in Table 3, reveal several key insights. In column (1), we estimate a coefficient of approximately 0.7 on $Middle_i \times Post_t$, akin to the result in Table 2. This coefficient represents the effect of the reduction on the middle class when advertising spending is zero. The marginal effect of advertising on the middle class after the reduction, captured by the triple interaction $Middle_i \times Post_t \times Adv_{s,t}$, is statistically insignificant. We find similar results in columns (2)-(3) after introducing additional interactions between the controls from Table 2 and our key advertising measures ($Post_t$, $Adv_{s,t}$, and $Post_t \times Adv_{s,t}$), which allows for targeted advertising by characteristics like income and age that correlate with wealth. To be clear, these results do not imply that the advisor's advertising had no impact on overall robo participation; this overall impact would be subsumed by the state-by-time fixed effects. Rather, the evidence implies that shifts in advertising cannot explain the *relative* increase in middle-class participation after the reduction.

[20] This could be attributed to the dual goals of expanding its long-term client base and enhancing short-term revenues. Notably, the bulk of short-term revenues primarily stems from wealthier clients, as the advisor only charges a fee for portfolios exceeding $10,000.

Lastly, as an additional robustness check, we introduce a different measure of state-level ad spending, denoted $Adv_{w,s,t}$, that accounts for targeting by market areas within states. This measure calculates the average spending in state s in the month of week t, separately for middle-class and upper-class DMAs. Since this measure varies within state-months, we extend our previous specifications from columns (1)-(3) by controlling for $Adv_{w,s,t}$ and $Post_t \times Adv_{w,s,t}$ separately. The results in columns (4)-(6) mirror those obtained from the initial measure.[21] We therefore conclude that the estimated effect of the reduction does not confound shifts in the advisor's advertising.

5.3 Additional Extensions

Gambling. Experimental evidence shows that investors exhibit lower risk aversion in the context of small lotteries (e.g., Bombardini and Trebbi 2012). So, lowering the minimum may induce some account holders to set up an additional, smaller "play account" in which to gamble. Whether this gambling channel drives the results in Table 2 depends on the relative size of a gambling account for the middle versus upper classes. If the optimal size of a gambling account for the upper class lies between \$500 and \$5,000 while the optimal size for the middle class lies below \$500, then Table 2 actually underestimates the effect of relaxing minimum-account constraints. If, on the other hand, the optimal size of a gambling account for the middle class lies between \$500 and \$5,000 while the optimal size for the upper class lies above \$5,000, then the gambling channel may actually drive the main results in Table 2. We assess this possibility by re-estimating the main specifications from Table 2 after removing the 12% of account holders with multiple accounts from the sample. The resulting estimates in Appendix Table A6 are almost the same as with the full sample, suggesting that gambling motives indeed do not drive the results.

Business Stealing. We examine whether our results are driven by a reallocation of actual or potential participants across robo advisors ("business stealing"), versus an overall increase in robo participation. A close examination of these concerns requires detailed data on clients of competitors. While we do not have access to such data, we nonetheless propose several basic tests based on publicly available data. Specifically, we use data from the SEC's Form ADV to examine trends in participation at other standalone robo advisors, namely Betterment and Personal Capital. The Form ADV data have limitations described in Section 3, but they are the best source of data

[21]The slight drop in the number of observations, less than 10%, between columns (1)-(3) and (4)-(6) can be attributed to certain states where all DMAs are grouped into a single wealth category. For instance, in Alabama, all DMAs (Birmingham, Huntsville, Mobile, and Montgomery) are classified as middle-class DMAs. As a result, $Adv_{w,s,t}$ is not defined for all upper-class households in Alabama within the deposit dataset. Consequently, these specific observations are excluded from the estimation in columns (4)-(6).

for this exercise, short of having microdata from each major U.S. robo advisor.

Business stealing could affect our results in two key ways. First, the reduction in the account minimum at the advisor may attract *existing* clients from competitor robo advisors. In this case, we would expect to observe a decline in the total number of clients at the competitors after the reduction. The results in Figure 8 do not support this concern; we actually observe an increase in the number of participants at competitors from the pre-reduction to the post-reduction periods.

Second, the reduction may attract *potential new* clients. In particular, 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 the advisor instead. This concern would imply that the competitors continue growing after the reduction, but at a lower rate. The results in Figure 8 are inconsistent with this concern since they show very little change in growth trends for competitors. We observe an increase in client growth only for the advisor itself, which is consistent with the effects of the reduction.

6 Description of Model

At an intuitive level, our reduced-form results may seem unsurprising. However, it is not obvious why retail investors would want a "textbook" robo portfolio given recent evidence that, in many cases, they prefer FinTech products that perpetuate their behavioral biases (e.g., Cookson, Engelberg and Mullins (2023); Barber et al. (2022); Ben-David et al. (2022)). This contrast raises the question of whether robo participants act in a welfare-improving manner. Moreover, even if the decision to become a robo investor increases welfare, the drivers and the distribution of any welfare gains are unknown. Studying welfare along these various dimensions lets us examine whether and how robo portfolios differ from existing financial products (e.g., index funds, TDFs), which informs why the robo market has continued to grow despite the availability of potential alternatives. We pursue these questions by calibrating a benchmark model of portfolio choice, suitably modified for our setting.

6.1 Setup

We follow the setup of benchmark portfolio choice models as closely as possible. See Gomes (2020) for a summary. We emphasize the unique features of our model below, while Appendix C provides full details on the more conventional aspects.

6.1.1 Preferences and Labor Income

Let i index investor. Time is discrete, and t indexes year. For the rest of the exposition, we conserve notation by aligning an investor's age with the year, such that we do not maintain both age and time subscripts. Investors begin their problem at age t_0. With probability p_t, an investor of age t survives until age $t+1$, and at age $\overline{T}$ any surviving investors leave the model. Investors have CRRA preferences over annual flow consumption $C_{i,t}$. Expected lifetime utility is given by:

$$U_{i,t} = \mathbb{E}_t \left[\sum_{\tau=t}^{\tau=\overline{T}} \delta^{\tau-t} \left(\prod_{j=t}^{j=\tau-1} p_j \right) \frac{{C_{i,\tau}}^{1-\gamma}}{1-\gamma} \right], \tag{6}$$

where δ is the discount factor, and γ is the coefficient of relative risk aversion. Investors enter age t with consumable and investable resources $W_{i,t}$. To match our empirical work, we call $W_{i,t}$ "liquid assets" instead of "cash-on-hand" (e.g., Deaton (1991)). Liquid assets are replenished through labor income and income from financial assets.

Appendix C.1.1 describes the labor income process, which is standard. Briefly, investors receive uninsurable labor income for $t \leq \underline{T}$, and retire at age $\underline{T}$+1. Following Carroll (1997), log income evolves according to the sum of: a deterministic function of age, f_i ; a permanent shock with volatility σ_v; a transitory shock with volatility of σ_ξ; a disaster shock that yields zero labor income and occurs with probability ϕ; and an aggregate shock that has a correlation coefficient ρ^Y with the robo portfolio's abnormal return, so that ρ^Y captures the cyclicality of labor income.

Since our empirical setting concerns pre-retirement investment, we err on the side of parsimony and avoid additional parameters that govern investment decisions within retirement. Instead, retired investors simply solve a consumption and savings problem with initial wealth $W_{\underline{T}+1}$. This setup abstracts from two other forms of retirement income: payments from an employer-sponsored plan (e.g., pension, 401k); and social security payments. Abstracting from the former seems reasonable, since only 26% of investors in the bottom three wealth quintiles have retirement support from their employer according to the 2016 SCF. Abstracting from the latter has a theoretically ambiguous effect that depends on the design of the social security system, a complicated problem that deserves its own study.[22]

[22] Following the modelling decisions of early papers (e.g., Campbell et al. (2001)), incorporating social security would function similarly to incorporating forced savings in a riskless asset. This would reduce investors' after-tax savings rate but increase their risky share within after-tax savings. The first channel raises the wealth threshold at which investors become a robo participant, which would shift the welfare gains from reducing the account minimum towards more-wealthy segments of the middle class. The second channel has the opposite effect. Our model overshoots the increase in robo participation by the second U.S. wealth quintile (Figure 9), which would loosely suggest that the channel related to a lower savings rate dominates.

6.1.2 Financial Assets

In modeling financial assets, we make two principal additions to benchmark models. First, rather than investing in a single, perfectly diversified risky asset, investors have two risky investment opportunities: a self-managed portfolio ($\mathcal{S}$), which can be likened to a discount brokerage account; and a portfolio overseen by a robo advisor ($\mathcal{A}$). We do not make any prior assumptions on the characteristics of these portfolios, instead letting the data inform these characteristics. Second, the two portfolios differ in that the robo portfolio requires an account minimum, M. We simplify the model's computational complexity by assuming investors cannot hold the self-managed and automated portfolios concurrently. This simplification has little bearing on the results because it is rarely optimal for investors to hold both portfolios at the same time.

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

$$R_{i,t}^{\mathcal{P}} = \beta_i^{\mathcal{P}} F_t + \epsilon_{i,t}^{\mathcal{P}}, \tag{7}$$

where F_t is a vector of priced risk factors, normally distributed with mean π^F and covariance matrix Σ^F; $\beta_i^{\mathcal{P}}$ is the loading of portfolio $\mathcal{P}$ on F_t for investor i; and $\epsilon_{i,t}^{\mathcal{P}}$ is an idiosyncratic shock, normally distributed with mean zero and volatility of $\sigma_{\epsilon,i}^{\mathcal{P}}$. The quantities of compensated, systematic risk ($\beta_i^{\mathcal{P}}$) and idiosyncratic risk ($\sigma_{\epsilon,i}^{\mathcal{P}}$) also vary across investors. This flexibility allows us to capture the personalization of robo portfolios, documented in Figure 1.

6.1.3 Model Solution

Appendix C.1.2 states the consolidated problem, as a Bellman equation. The state variables are age (t) and liquid assets ($W_{i,t}$). The control variables are consumption and the share of liquid assets allocated to: the risk-free asset, the self-managed portfolio, and the robo portfolio. We solve the model through backward induction, using standard numerical methods in Appendix C.1.3.

6.2 Calibration

6.2.1 Portfolio Parameters

We next discuss the calibration of the portfolio parameters $\left\{\sigma^{\mathcal{S}}_{\epsilon,i}, \sigma^{\mathcal{A}}_{\epsilon,i}, \beta^{\mathcal{S}}_{i}, \beta^{\mathcal{A}}_{i}\right\}$. We use the portfolio dataset, which contains pairs of self-managed and robo portfolios (see Section 3.2). Our calibration proceeds in three steps, briefly summarized below. Appendices C and D have details.

First, following a convention in the household finance literature (e.g., Calvet, Campbell and Sodini (2007); Von Gaudecker (2015)), we estimate factor loadings and idiosyncratic volatilities for all individual securities. Specifically, for a given vector of risk factors, F, we estimate the following pricing equation for each security k in the portfolio dataset,

$$R_{k,m} = \beta_k F_m + \epsilon_{k,m}, \tag{8}$$

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 monthly percent changes in the U.S. and the global bond indices from Barclays.[23]

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 pair of portfolios, indexed by j. Let $w^{\mathcal{S}}_{j}$ denote a vector of weights across securities k for the self-managed portfolio j, and let $w^{\mathcal{A}}_{j}$ denote the weight vector for j's matched robo portfolio. Given an estimated vector of loadings across securities, $\hat{\beta}$, and covariance matrix of idiosyncratic volatilities, $\hat{\Sigma}_{\epsilon}$, we calculate the portfolio parameters as

$$\sigma^{\mathcal{P}}_{\epsilon,j} = \sqrt{w^{\mathcal{P}\prime}_{j}\hat{\Sigma}_{\epsilon}w^{\mathcal{P}}_{j}}, \;\; \beta^{\mathcal{P}}_{j} = w^{\mathcal{P}\prime}_{j}\hat{\beta}, \tag{9}$$

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

Table 4 compares benchmark characteristics of self-managed and robo portfolios. Robo portfolios feature a 30 pps higher Sharpe ratio for both the middle and upper classes. The higher Sharpe ratio reflects a 2 pps higher expected return and a 11 pps lower idiosyncratic volatility. By construction, the higher expected return stems from greater exposure to priced risk factors, specifically value and bond premia (Appendix Table A13). Bonds, in particular, have a relatively high Sharpe ratio, which undergirds the well-known strategy of risk parity (e.g., Asness, Frazzini and Pedersen (2012)). The lower idiosyncratic volatility in robo portfolios reflects how investors

[23] We calibrate the mean and covariance matrix of F using the longest observed time series over 1960-2017 and report these values in Appendix Table A14. For reference, the mean and volatility of the market factor equal 7.6% and 14.7%, respectively. Our baseline choice of five factors includes many of the factors to which the robo advisor gives exposure. However, we obtain similar insights from other factor models (Appendix Table A15).

in our data mostly hold individual stocks and actively-managed funds. For example, the median self-managed portfolio allocates only 8% to broad-based index-linked ETFs.

To understand the magnitude of the improvement, we perform a return-loss calculation similar to Calvet, Campbell and Sodini (2007). This calculation gives the counterfactual expected return that the investor would receive if they invested with the same total volatility, but with the greater Sharpe ratio of its robo match. We call this counterfactual return the "robo-Sharpe expected return", and the difference relative to investors' actual return gives the return loss from investing inefficiently on their own. Accordingly, the average investor in our sample would gain 7 pps in expected return from investing with the same efficiency as the robo advisor. This potential gain is the same for both wealth classes. Appendix Table A11 recalculates all of these results without the approximation that the advisor can hold fractional shares for portfolios smaller than \$5,000. The results are similar because, as described in Section 2.2, the advisor substitutes the primary ETF with a highly-correlated secondary ETF that has a lower share price if holding the primary ETF in discrete quantities would induce meaningful tracking error. Appendix Table A12 verifies that the results are robust to excluding self-managed accounts smaller than 10% of the investor's wealth, which may be gambling accounts.

Lastly, in the third step of the calibration, 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 investor'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 investor i.

6.2.2 Other Parameters

Table 5 reports the other parameters. We choose standard preference parameter values of $\gamma = 4$ and $\delta = 0.95$. Appendix Table A19 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 investor 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 $\rho^Y = 0.3\%$. Following Carroll (1997), we parameterize $\phi = 0.12\%$ to match the share of investors 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)).

6.3 Characterizing Investor Behavior through Policy Rules

The model reproduces many common implications of life cycle models, but it also generates several implications that arise due to our unique setting. We succinctly summarize investors' optimal investing behavior (policy rules) in Figure 8. Appendix C.2 has more detailed discussion.

The following patterns are conventional. Young investors accumulate a buffer stock of liquid assets to insure against idiosyncratic income shocks. Consequently, the savings rate initially decreases in age (panel (a)). As investors approach retirement, this pattern reverses and the savings rate begins to increase. The youngest investors also hold most 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 investors in their 40s (panel (b)). Lastly, since we do not incorporate stock market participation costs until Section 7.4.2, almost all investors hold risky assets, either through the robo portfolio or through the self-managed portfolio.

Unlike in benchmark models, in our model investors choose between two portfolios. On the one hand, the robo portfolio is preferable because of the advantages just described (e.g., diversification, priced risk exposure). On the other hand, the robo portfolio has an account minimum. Due to this trade-off, investors switch to the robo portfolio after their liquid assets cross some threshold. Panel (d) shows that, under a \$5,000 account minimum, the threshold ranges between \$25,000 and \$30,000 for most of the investor's working life. This result is broadly consistent with the data, as the median liquid assets of existing middle class robo participants equals \$25,000 (Table 1). Once the investor becomes a robo participant, the risky share jumps (panel (b)) due to the advantages of robo portfolios.

As investors approach retirement, the threshold falls for two reasons. First, as in benchmark models, older investors have a higher savings rate. So, they are willing to allocate a large share of their resources to the robo advisor. Second, unlike in benchmark models, robo portfolios feature a double glide path, taking less priced risk as the investor ages or suffers negative wealth shocks. This personalization strengthens older investors' demand for the robo portfolio, and, by extension, raises their risky share. This channel offsets the decrease in risky share that would obtain in a benchmark model from the declining share of human capital in total wealth.

7 Effect of the Reduction within the Model

7.1 Can Portfolio Differences Explain the Empirical Effect?

We compute the reduction's theoretical effects as follows. We first use the policy functions to calculate how an investor of a given age and level of wealth optimally invests under the minimums of \$5,000 and \$500. We then aggregate these individual decisions using the joint distribution of age and wealth in the 2016 SCF. Precisely, we discretize this joint distribution, and we average investment decisions across age-by-wealth cells weighting by the share of the overall U.S. population in each cell. This approach differs from simulation-based methods that would instead rely on a model-implied wealth distribution, obtained from simulating an investor from starting age to retirement. This distinction matters, since in simulated models investors rapidly accumulate wealth as they age (e.g., Gomes (2020)). Consequently, the reduction would only affect the youngest investors, as most older investors would have already overcome the \$5,000 minimum. In reality, however, investors accumulate wealth much more slowly. For example, nearly 40% of U.S. retail investors over age 40 have less than \$6,000 in liquid assets (Appendix Figure A4). Our methodology accounts for this empirical regularity.

Even though we begin with the external wealth distribution, the model will still generally undershoot the level of welfare gains because the gains are calculated using internal policy rules. Since investors save aggressively under such policy rules, those who become robo participants in response to the reduction would have otherwise overcome the \$5,000 minimum in a few years. Consequently, they value access to the robo advisor at the lower \$500 minimum less than they would under the more-realistic scenario in which they slowly accumulate wealth.

7.1.1 Results

We focus on comparing the reduction's theoretical and empirical effects on the wealth distribution of robo participants. In Figure 9, we overlay the empirical effect from Figure 3 with red open circles that show the theoretical effect. The model matches the share of existing participants from each wealth quintile almost exactly. It also explains the wealth distribution of new participants quite well. This result is surprising because of the model's parsimony. Given that our goal is to ask how far a benchmark model can go in explaining the data, we refrain from adding parameters that would fine-tune the model's fit but cannot be quantitatively disciplined in a straightforward way (e.g., fixed costs of delegation).

The model's consistency with the data also indirectly supports our main DiD identification assumption. In the model, the difference in the reduction's effect on the middle versus the upper class reflects the effect of relaxing minimum-account constraints, by construction. This is exactly what we assume in our DiD research design, thereby supporting the validity of the estimates.

To gain more confidence in the model, we examine whether it can match other features of the data. Table 6 evaluates the model's fit through three other sets of statistics. Panel (a) reports the pre-reduction and post-reduction share of investors from the middle-class. This share is essentially a weighted average of the second and third quintiles shown in Figure 9. Panel (b) reports the risky share for new and existing middle-class participants. The model matches the data well on average, though producing a lower risky share for new middle-class participants. Lastly, panel (c) presents the model-implied stock market participation rate, after adding standard stock market participation costs (Section 7.4.2). Interestingly, the model matches the U.S. stock market participation rate very well.

These findings suggest that middle-class investors seek robo advisors largely because of the portfolio-level advantages that they confer. In particular, these advantages alone can mostly explain the empirical effects of the reduction. This interpretation accords with surveys, which report improved asset allocation as a major reason why retail investors seek robo advisors (Rossi and Utkus (2021a), Costa and Henshaw (2022)). We return to the question of why investors cannot replicate the robo portfolio on their own in our conclusion.

7.2 Welfare Gains

We move to the model's second purpose: computing the welfare gains from access to robo advisors. We are particularly interested in the channels that drive any welfare gains and the distribution of these gains. As standard, we measure investor i's welfare gain by the percent increase in annual consumption under the previous minimum that raises their expected lifetime utility by the same amount as the reduction. Let $\underline{V}_i$ and $\overline{V}_i$ denote the value functions under the minimums of \$5,000 and \$500, respectively. Appendix C shows that this welfare gain equals:

$$q_i = \left(\frac{\overline{V}_i}{\underline{V}_i}\right)^{\frac{1}{1-\gamma}} - 1. \tag{10}$$

Table 7 reports various calculations of equation (10). Panel (a) reports the welfare gain under various specifications of the self-managed portfolio. We focus primarily on the first row of panel (a), in which self-managed portfolios are calibrated using investors' actual allocations in the portfolio dataset as described in Section 6.2.1. Column (1) shows that the baseline welfare gain equals 0.77% of lifetime consumption, which is moderate by the standards of the literature, per footnote 4. For reference, this gain lies within the range of gains from correcting various investment mistakes in the analogous specification of Cocco, Gomes and Maenhout (2005), such as ignoring labor income (0.65%) or not investing in risky assets (1.04%).

7.2.1 Human Capital and Heterogeneous Welfare Gains by Age

We next examine how welfare gains vary by age. Returning the benchmark comparison group based on actual self-managed portfolios, the results in columns (2)-(4) show that new participants older than 55 gain almost three times as much (1.68%) as new participants aged 35 or younger (0.58%). This finding challenges robo advisors' popular image and their claim to "build our products and services for millennials" (e.g., Hutchins (2020)).

The central explanation for this difference lies in human capital. Older investors have fewer remaining working years, as well as non-increasing earnings growth. Consequently, they may fail to reach the wealth threshold at which it becomes optimal to participate with the robo advisor before retirement, under a \$5,000 minimum. By contrast, younger investors may eventually accumulate enough earnings to become robo investors even under this higher minimum. As a result, older investors benefit more from the reduction than their younger peers. This argument is especially relevant in the context of the overall U.S. population, in which investors with modest wealth comprise a large share of investors past their peak earnings years (Appendix Figure A4).

We sharpen this explanation by asking how many years it would take for a robo non-participant to accumulate enough wealth to become a robo participant under a minimum of \$5,000. We simulate 5,000 life cycles of robo participation from various starting ages. We next calculate the average number of years until the investor's first robo investment across simulations. We normalize this average by dividing it by the remaining years in the investor's working life.[24]

Figure 10 shows that a 55 year-old who does not participate under a \$5,000 minimum would continue to be a robo non-participant for around half of her remaining working life. By contrast, a similar 25 year-old would spend only around 3% of her life as a non-participant. This result corroborates the previous intuition: the 55 year-old has stronger demand for robo portfolios because

[24]The starting level of liquid assets equals \$6,000 in each simulation, which is the 40$^{\text{th}}$percentile of the U.S. wealth distribution (i.e., the midpoint of the middle class). The starting level of permanent income $\nu_{i,t}$ equals the average value among investors who have accumulated less than \$6,000 by the time they reach the starting age.

she is much less likely to accumulate enough earnings to overcome the higher minimum.[25]

7.3 Portfolio Characteristics that Drive the Results

By construction, any welfare improvement in the model must reflect a difference between robo and self-managed portfolio characteristics. We next assess each characteristic's welfare impact to infer the sources of middle-class investors' demand for wealth management.

7.3.1 Inference through Counterfactual Robo Portfolios

Our most direct approach to this question calculates the welfare gain from reducing the minimum on a counterfactual robo portfolio that no longer differs from the self-managed portfolio in terms of a given characteristic. The difference between the counterfactual gain and the baseline gain in panel (a) tells us how important that characteristic is from a welfare perspective.[26]

We report the results in panel (b) of Table 7, starting with differences in exposure to idiosyncratic risk ($\sigma_{\epsilon,i}^{\mathcal{P}}$) and to priced risk ($\beta_i^{\mathcal{P}}$). Column (1) shows a substantially smaller welfare gain of 0.23% if the diversification benefit of robo portfolios is stripped away (i.e., no reduction in $\sigma_{\epsilon,i}^{\mathcal{P}}$). Similarly, we obtain a gain of 0.2% if we remove the difference in robo portfolios' bond exposure. This reflects the relatively high Sharpe ratio of bonds discussed in Section 6.2.1. By contrast, constraining robo portfolios to provide the same stock market beta as self-managed portfolios results in almost the same welfare gain (0.72%). This finding reflects how self-managed and robo portfolios already have similar stock market exposures. Lastly, removing the difference in value (HML) and size (SMB) exposures results in a smaller gain of 0.46%. This gain arises because self-managed portfolios have less exposure to value stocks (Appendix Table A13).

We next consider the welfare effects of the double glide path. 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 to nei-

[25]This finding does not contradict column (3) of Table 2 because it shows that the magnitude of the welfare gain varies by age, whereas Table 2 shows that the effect on participation does not vary by age. In particular, any investor with a positive welfare gain will choose to participate. So, our model interprets Table 2 as saying that the probability of experiencing a positive gain does not vary by age, even though older investors gain more overall. Consistent with this interpretation, we replicate column (3) of Table 2 in our model and find that the average change in the probability of robo participation for investors under age 50 (+20 pps) is almost the same as for investors above age 50 (+19 pps).

[26]An alternative approach would be to remove differences in characteristics one after another. However, this approach might understate or overstate the gains from each characteristic. For example, if we first remove the difference in idiosyncratic risk, the gain equals 0.23% (a decline of 0.54 pps). By construction, removing any other characteristic after this will not reduce the gain by more than 0.23%. Based on this result, one may conclude that the difference in idiosyncratic risk is the most important. However, if we first remove the difference in bond loadings, the gain is an even smaller 0.20% (a decline of 0.57 pps). The order of removals matters because the consumption metric is a nonlinear function of the welfare change, in utils, and because there are complementarities between different characteristics. Therefore, we undertake a more conservative approach, removing one characteristic at a time.

ther vary by age nor by wealth. Panel (b.ii) reports the welfare gain from reducing the minimum on these counterfactual robo portfolios. We find equally smaller gains when removing the adjustment for age (0.34%) or the adjustment for wealth (0.37%). This result suggests that middle-class investors value both adjustments comparably. Column (4) shows that the adjustment by age is particularly important for older investors.[27]

Our counterfactual analysis suggests that middle-class investors value both basic and sophisticated characteristics of robo portfolios. At a basic level, robo portfolios add value through improving diversification. On its own, this finding may suggest that robo advisors may offer similar benefits as buying an index fund. However, investors also place significant value on arguably more-sophisticated characteristics, like exposure to additional stock and bond risk factors and a double glide path by age and wealth. Combining these characteristics with improved diversification is important for generating welfare gains, as we shall see from the next approach that directly computes the welfare gain for someone who already invests in an index fund.

7.3.2 Inference through Hypothetical Comparison Groups

The previous approach asks how much actual investors would gain when reducing the minimum on a counterfactual robo portfolio. Symmetrically, one can also ask how much a counterfactual investor would gain when again reducing the minimum on an actual robo portfolio. To do so, we reperform the baseline welfare calculation in panel (a) of Table 7 after replacing the self-managed portfolio with an alternative comparison. Motivated by the results in panel (b), we specify two comparisons: a stock index fund, specifically Vanguard's Total Stock Market ETF (VTI); and the Vanguard TDF with a target year that an investor seeking to retire at age 65 would have chosen in 2015.

Panel (c) of Table 7 summarizes the results. A new middle-class participant whose risky portfolio would consist only of VTI without the reduction gains 0.72% in lifetime consumption, or around 90% of the gain for one who would otherwise invest as do the investors in our data. So, while robo portfolios partly resemble a stock index fund, as suggested by Reher and Sun (2019),

[27] Personalization by wealth 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)). Intuitively, labor income functions like an endowed risky asset. Therefore, the optimal asset allocation, as encoded by the loadings $\beta_i^{\mathcal{P}}$, will generically depend on the share of income in total wealth, which, in turn, depends on both the investor's age and wealth. Additionally, the two-fund theorem does not necessarily hold because the account minimum and the implicit borrowing constraint place bounds on the risky share. Lastly, while each household has an optimal robo allocation, the actual robo allocation only approximates the optimal one. Personalization raises welfare insofar as it moves a household closer to her optimal allocation.

investors appear to value characteristics of robo portfolios that go beyond this similarity. Otherwise, this hypothetical investor would have experienced no gain.[28]

Next, consider a hypothetical middle-class participant whose risky portfolio would have consisted of a Vanguard TDF without the reduction. To focus the comparison on asset allocation, we ignore the fact that this TDF itself has an account minimum of $1,000, which means we provide a lower bound on the welfare gain. Accordingly, this hypothetical investor gains 0.23% in lifetime consumption, or one-third of the gain for a participant who would have otherwise invested as investors actually do in our data.[29] On the one hand, both TDFs and robo portfolios go beyond a stock index fund by giving investors access to bond risk factors and by adjusting their allocation as they age, as documented in Appendix Figure A7. This similarity explains why we find a smaller gain for a hypothetical investor who already invests in a TDF. On the other hand, robo portfolios also adjust for changes in the investor's wealth, which explains why the gains are still non-trivially positive. Quantitatively, the results make sense in light of the companion exercise in panel (b), which finds gains of the same magnitude when removing the distinguishing features of TDFs from robo portfolios (bond exposure and age glide paths).

Based on our findings from this subsection, we conclude that the recent growth in the robo market differs from the concurrent growth observed in index funds and TDFs. While these trends share a connection, the distinction lies in the fact that investors would still gain from access to robo advisors even if they already invested in one of these alternative products. This is because robo portfolios introduce additional layers of sophistication beyond the features already in TDFs and, especially, in index funds. Hence, even a hypothetical investor with a relatively sophisticated approach of investing in TDFs would nevertheless benefit from accessing robo portfolios. In a broader context, this finding aligns with the advisor's mantra that "everyone deserves sophisticated investment advice", emphasizing that a significant portion of the gains from accessing robo portfolios stems from their more advanced characteristics.

[28]This finding does not contradict the large gains from diversification implied by panel (b.i) because the efficient benchmark in our model is not necessarily a stock index fund. In particular, returns do not follow the CAPM but rather the factor model in equation (8), which accounts for multiple stock and bond risk factors. This also explains why the distribution of welfare gains across age is more pronounced relative to panel (a). Specifically, investors in our data do have some exposure to bond factors. Since tilting towards bonds and away from stocks is optimal for older investors (Figure 8), the actual older investors in our data gain less from the reduction than their hypothetical peer who would otherwise invest only in VTI.

[29]We obtain a similar gain of 0.16% when using the Fidelity Freedom Fund TDF as the comparison (Appendix Table A18). Most TDF target years vary in 5 year increments, and so we assign investors to the TDF that would deliver a retirement age between 65 and 69 based on their age in 2015.

7.4 Additional Analysis

7.4.1 Heterogeneity by Pre-Participation, Initial Deposit, and Wealth

Our main welfare results in Table 7 restrict to new robo participants, but it is possible that existing robo participants also gain from the reduction. Indeed, over half of existing middle-class participants invested precisely \$5,000 as their initial deposit, per Table 1. Such investors may have wanted to invest less than \$5,000 were it not for the account minimum, but they nevertheless prefer investing precisely \$5,000 to not investing at all. Reducing the account minimum to \$500 enables such investors to reduce their robo allocation to its optimal level. Panel (a) of Appendix Table A16 uses the model to calculate the welfare gain for all middle-class robo participants, including existing ones who experience positive gains. The overall gain equals 0.62%, which is a weighted average of the 0.77% gain for new participants already shown in Table 7 and a much smaller 0.05% gain for existing participants. Thus, the intensive margin effect on existing participants attenuates the extensive margin effect on new participants, leading to a smaller welfare gain for middle-class participants overall than for new middle-class participants.[30]

Panel (b) of Appendix Table A16 again restricts to new robo participants, and it separately calculates the welfare gains for investors whose initial deposit lies in various ranges below \$5,000. We find gains of 0.86%, 0.70%, and 0.55% for investors with an initial deposit in the ranges from \$500 to \$2,000, from \$2,000 to \$3,500, and from \$3,500 to \$4,500, respectively. Panel (c) performs a similar exercise in terms of the investor's wealth quintile, finding a larger gain of 1.21% for new robo participants in the second U.S. wealth quintile versus 0.52% for those in the third quintile. The economic intuition for these findings reflects how the optimal robo investment is increasing in the investor's wealth (Figure 8).[31] Without the reduction, therefore, new participants from the

[30] An analogy can be drawn between this finding and the analysis of extensive and intensive margin effects in Reher and Sun (2019). Specifically, Reher and Sun (2019) demonstrate that reducing the account minimum increases log participant inflow from less-wealthy investors by twice as much as it raises log dollar inflow. In their context, extensive versus intensive margins pertain to the number of accounts versus dollars-per-account: the reduction raises the number of less-wealthy participants (extensive margin); however, due to their small account size (intensive margin), the effect on participant inflow surpasses the effect on deposit inflow. In our contribution, our focus lies on understanding the drivers and distribution of welfare. Therefore, the concept of extensive versus intensive margins pertains to welfare gains for new versus existing robo participants. Specifically, Appendix Table A16 suggests a significant effect on new robo participants (extensive margin) but a minor effect on existing ones (intensive margin), leading to a more muted overall gain. Just as the intensive margin in Reher and Sun (2019) attenuates the reduction's effect on overall dollar inflow from their definition of the less-wealthy, the intensive margin in our setting attenuates the reduction's impact on overall welfare gains for the middle class.

[31] Returning to the data, Appendix Table A17 substantiates this point by re-tabulating several of the summary statistics from Table 1 according to the investor's initial deposit and wealth quintile. Consistent with the model, the results in panels (a)-(b) of Appendix Table A17 imply that new participants from the second wealth quintile make smaller initial deposits relative to those from the third quintile, and also relative to existing participants from the second quintile.

second quintile would have required a longer time to accumulate enough wealth to be willing to invest, relative to new participants from the third quintile. Consequently, they value the reduction in account minimum much more.

Collectively, these findings imply that the reduction yields the largest gains for middle-class investors who respond by making very small initial deposits (e.g., under $2,000). However, their larger gains are diluted by smaller gains for either existing participants or new participants who do not make such a small deposit.

7.4.2 Model Extensions and Robustness

Appendix C.4 describes various extensions and robustness exercises, with results in Appendix Table A18. Briefly summarizing:

Participation Costs. We calculate a higher welfare gain of 1.23% when, as in Vissing-Jørgensen (2002), investors incur a per-period cost of holding the self-managed portfolio ($200). This gain is of the same order of magnitude as our baseline gain, reflecting how non-participation in the stock market leads to modest welfare losses if the alternative is investing inefficiently on one's own (e.g., Calvet, Campbell and Sodini (2007)).

Borrowing. Allowing investors to borrow at the interest rate on consumer credit has little impact on the results. Intuitively, borrowing to invest in the robo portfolio would be suboptimal given the high interest rate on consumer credit.

Alternative Calibrations. We find similar results when: recalibrating the portfolio parameters by using the parameters in Calvet, Campbell and Sodini (2007); and recalibrating the preference parameters γ and δ in various ways. The appendix also estimates these parameters via GMM, finding similar values as in our baseline parameterization ($\gamma = 5.1$, $\delta = 0.92$) and a gain of 1.19%.

Procyclical Income. Lastly, we parameterize $\rho^Y = 80\%$, such that labor income becomes more stock-like. Thus, contrary to the glide paths found among traditional TDFs, the optimal risky share increases in age (e.g., Benzoni, Collin-Dufresne and Goldstein (2007); Lynch and Tan (2011)). On the one hand, such a parameterization should theoretically lower the value of investing with the robo advisor, as doing so increases the investor's total risk through correlation between labor and financial income. On the other hand, this correlation reduces an investor's unconstrained-optimal risky share, which makes investing under a lower minimum especially valuable. Moreover, concerning glide paths, the advisor's double adjustment by age and wealth attenuates the loss from how the age glide path reduces priced risk over the life cycle. On net, these forces approximately cancel out, leading to almost the same results as in our baseline calculation.

8 Conclusions

We draw two conclusions. First, wealth management is not necessarily a luxury service, and many middle-class investors have substantial demand for it. Recent FinTech innovation, in the form of robo advisors, can accommodate that demand. Robo advisors do so by using new technologies to expand access to sophisticated investing, thus providing retail investors with advanced features such as double glide paths and exposure to multiple risk factors. We provide novel evidence that these relatively-sophisticated features are an important driver of demand for robo advisors, and, quantitatively, they are at least as important for explaining this demand as more-basic features, like diversified exposure to the stock market. Expanding access leads to moderate gains overall for investors, but these gains vary considerably and are largest for older, lower-middle-class investors. These results help explain why the market for automated wealth management has continued to grow, both independently and within the traditional financial sector, despite the existence of products like index funds and TDFs.

Second, our results support models of bounded rationality in which retail investors optimally seek professional management given limits on their ability to invest with the same level of sophistication as a professional manager. By gaining access to robo advisors, less-wealthy investors experience moderate welfare gains. We leave open the question of why households do not invest like a professional, 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* . 6.2.1

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* . 6

Bachas, P., Gertler, P., Higgins, S. and Seira, E.: 2020, How Debit Cards Enable the Poor to Save More, *The Journal of Finance* . 6

Bagliano, F. C., Fugazza, C. and Nicodano, G.: 2018, A Life-Cycle Model with Unemployment Traps. C.1.1

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

Barber, B. M., Huang, X., Odean, T. and Schwarz, C.: 2022, Attention Induced Trading and Returns: Evidence from Robinhood Users, *The Journal of Finance* . 5, 6

Barber, B. M. and Odean, T.: 2000, Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors, *The Journal of Finance* . A.3.4

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* . 6

Ben-David, I., Franzoni, F., Kim, B. and Moussawi, R.: 2022, Competition for Attention in the ETF Space, *Review of Financial Studies* . 5, 6

Benzoni, L., Collin-Dufresne, P. and Goldstein, R. S.: 2007, Portfolio Choice over the Life-Cycle when the Stock and Labor Markets Are Cointegrated, *The Journal of Finance* . 7.4.2

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.: 2018, Financial Literacy and Portfolio Dynamics, *Journal of Finance* . 8

Bianchi, M. and Briére, M.: 2022, Augmenting Investment Decisions with Robo-Advice. 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* . 5.3

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* . 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* . 2.2, 6.2.1, 6.2.1, 7.4.2, C.4.1, C.4.3, D, D, A18

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*. 22

Campbell, J. Y. and Viceira, L. M.: 2002, *Strategic Asset Allocation: Portfolio Choice for Long-term Investors*, Oxford University Press. 27

Carroll, C. D.: 1997, Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis, *The Quarterly Journal of Economics* . 6.1.1, 6.2.2, 5, C.1.1, C.1.1

Carroll, C. D. and Samwick, A. A.: 1997, The Nature of Precautionary Wealth, *Journal of Monetary Economics* . C.1.1

Catherine, S.: 2022, Countercyclical Labor Income Risk and Portfolio Choices over the Life-Cycle, *Review of Financial Studies* . C.1.1

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

Chodorow-Reich, G., Nenov, P. T. and Simsek, A.: 2021, Stock Market Wealth and the Real Economy: A Local Labor Market Approach, *American Economic Review* . 5.2, 6, A.3.3, A.3.4

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* . 4, 6.2.2, 7.2, 5, C.1.1, C.1.2, C.1.3

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

Cookson, A. J., Engelberg, J. E. and Mullins, W.: 2023, Echo Chambers, *Review of Financial Studies* . 5, 6

Costa-Dias, M. and O'Dea, C.: 2019, Dynamic Economics in Practice, *Institute for Fiscal Studies Working Paper* . C.1.2

Costa, P. and Henshaw, J. E.: 2022, Quantifying the Investor's View on the Value of Human and Robo-Advice, *Vanguard Research Report* . 7.1.1

D'Acunto, F., Prabhala, N. and Rossi, A. G.: 2019, The Promises and Pitfalls of Robo-Advising, *Review of Financial Studies* . 3, 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

Fama, E. F. and French, K. R.: 1993, Common Risk Factors in the Returns on Stocks and Bonds, *Journal of Financial Economics* . 2.2.1, 6.2.1, 8

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

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* . 6

Fuster, A., Plosser, M., Schnabl, P. and Vickery, J.: 2019, The Role of Technology in Mortgage Lending, *Review of Financial Studies* . 6

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

Gentzkow, M.: 2006, Television and Voter Turnout, *The Quarterly Journal of Economics* . 5.2, A.3.3

Gentzkow, M. and Shapiro, J. M.: 2008, Preschool Television Viewing and Adolescent Test Scores Historical Evidence from the Coleman Study, *The Quarterly Journal of Economics* . 5.2, A.3.3

Gomes, F.: 2020, Portfolio Choice Over the Life Cycle: A Survey, *Annual Review of Financial Economics* . 1, 7, 6.1, 7.1, C.2

Gordon, B. R., Jerath, K., Katona, Z., Narayanan, S., Shin, J. and Wilbur, K. C.: 2020, Inefficiencies in digital advertising markets, *Journal of Marketing* . 19, A.3.3

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* . C.1.1

Higgins, S.: 2022, Financial Technology Adoption: Network Externalities of Cashless Payments in Mexico. 6

Hong, C. Y., Lu, X. and Pan, J.: 2022, FinTech Adoption and Household Risk-Taking, *NBER Working Paper* . 6

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. 7.2.1

Hvide, H. K., Meling, T. G., Mogstad, M. and Vestad, O. L.: 2023, Broadband Internet and the Stock Market Investments of Individual Investors, *The Journal of Finance* . 1

Jurek, J. W. and Viceira, L. M.: 2011, Optimal Value and Growth Tilts in Long-Horizon Portfolios, *Review of Finance* . 2.2.1

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* . 5.2

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. 3, 1

Lucas, R. E.: 1987, *Models of Business Cycles*, Oxford University Press. 4

Lynch, A. W. and Tan, S.: 2011, Labor Income Dynamics at Business-Cycle Frequencies: Implications for Portfolio Choice, *Journal of Financial Economics* . 7.4.2

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

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

Nash, A.: 2015, It's Time to Kill the Monthly Fee for Small Accounts, *Technical report*. 12

Newey, W.: 2007, Course Materials for Nonlinear Econometric Analysis, *MIT OpenCourseWare* . C.5, A19

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.: 2019, On FinTech and Financial Inclusion, *NBER Working Paper* . 1

Reher, M. and Sun, C.: 2019, Automated Financial Management: Diversification and Clientele Effects, *Journal of Investment Management* . 1, 7.3.2, 30

Roll, R.: 1977, A critique of the asset pricing theory's tests Part I: On past and potential testability of the theory, *Journal of Financial Economics* . 35

Rossi, A. G. and Utkus, S. P.: 2021a, The needs and wants in financial advice: Human versus robo-advising, *Working Paper* . 1, 7.1.1

Rossi, A. and Utkus, S.: 2021b, Who Benefits from Robo-Advising: Evidence from Machine Learning. 3, 1

Roussanov, N. L., Ruan, H. and Wei, Y.: 2020, Marketing Mutual Funds, *Review of Financial Studies* . 5.2

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

Tauchen, G. and Hussey, R.: 1991, Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models, *Econometrica* . C.1.3

Tobin, J.: 1958, Liquidity Preference as Behavior Towards Risk, *Review of Economic Studies* **25**. 27

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

Vissing-Jørgensen, A.: 2002, Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structure, *NBER Working Paper* . 7.4.2, C.4.1

Von Gaudecker, H.-M.: 2015, How Does Household Portfolio Diversification Vary with Financial Literacy and Financial Advice?, *The Journal of Finance* . 8, 6.2.1

Wealthfront: 2023, Automated Investing Accounts below $5,000, *Technical report*. 10, A11

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.3.5

Figure 1: Understanding the Robo Advisor's Asset Allocation

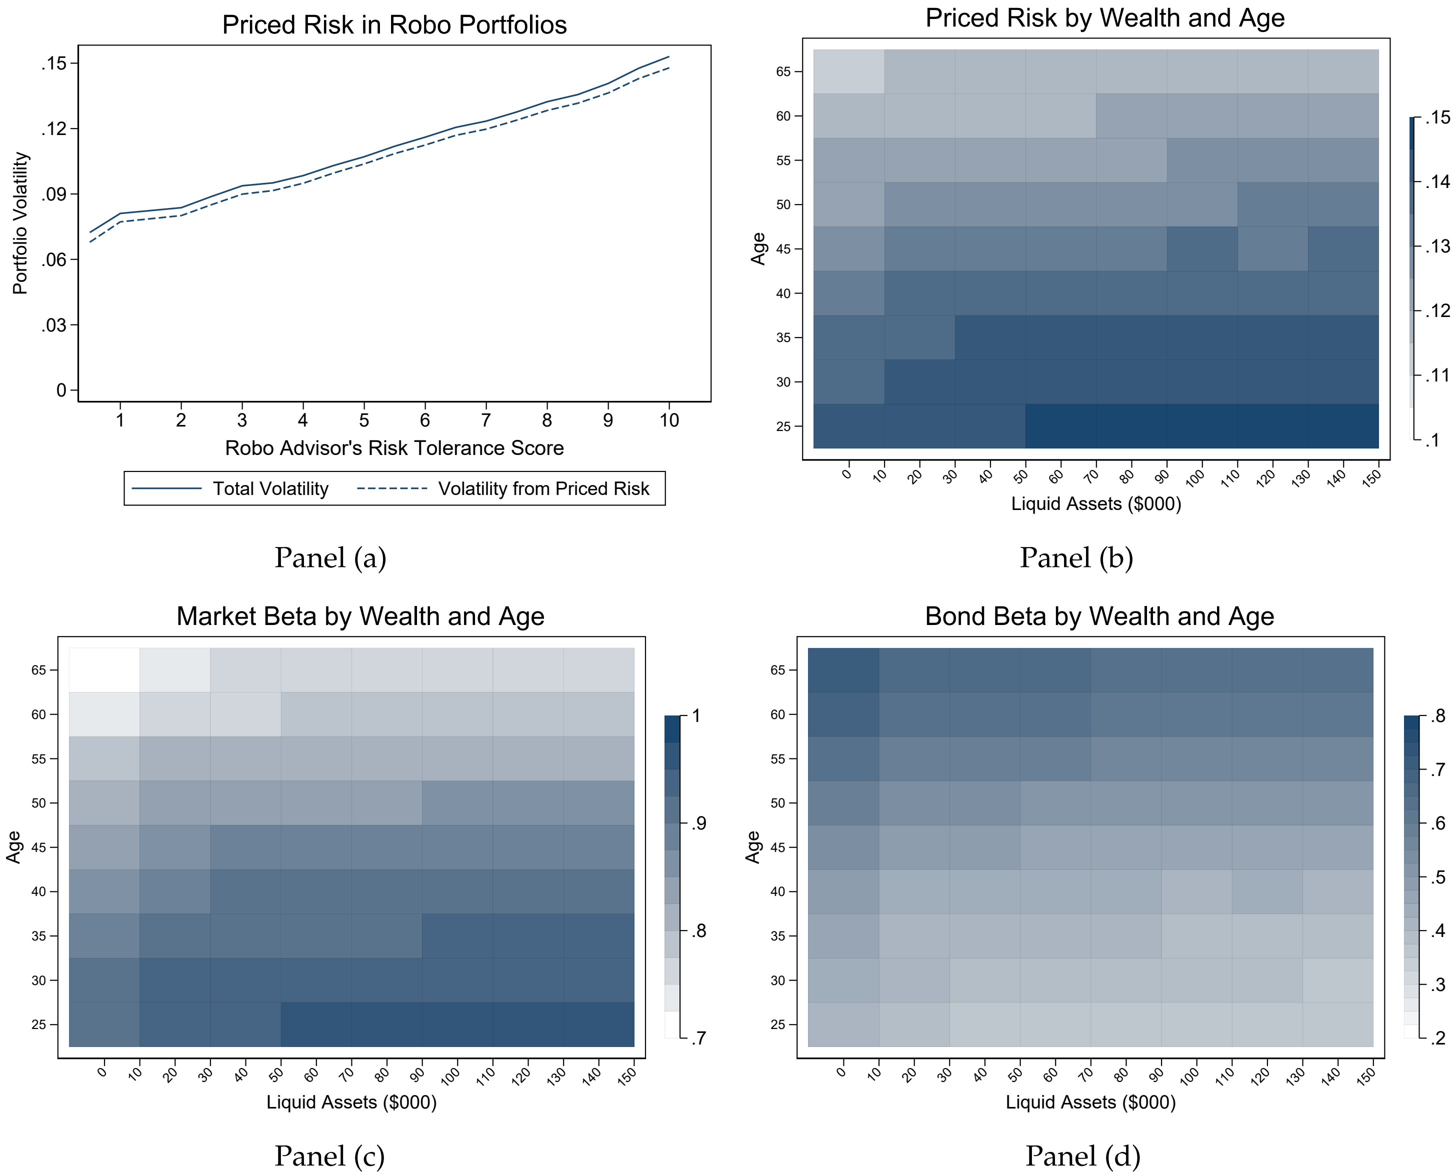

Note: This figure summarizes the priced risk from factor exposure in robo portfolios. Panel (a) plots the total volatility and the volatility from priced risk by the portfolio's risk tolerance score, which the robo advisor assigns based on the investor's response to the advisor's questionnaire. *Total Volatility* is the standard deviation of the total portfolio returns, which is the sum of idiosyncratic return and systematic return. *Volatility from Priced Risk* is the standard deviation of systematic return. Systematic returns are based on 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 investors of a given age and liquid wealth. The five factors are: the three Fama and French (1993) factors; and the monthly percent changes in Barclays' U.S. and global bond indices. Panel (b) plots volatility from priced risk by investor age and wealth. Panels (c)-(d) plot the loadings on the market factor (*Market Beta*) and U.S. bond factor (*Bond Beta*) by investor age and wealth. Details on robo portfolios are in Section 3. Details on estimating the factor models are in Appendix D.

Figure 2: How the Robo Wealth Distribution Changes after Expanding Access

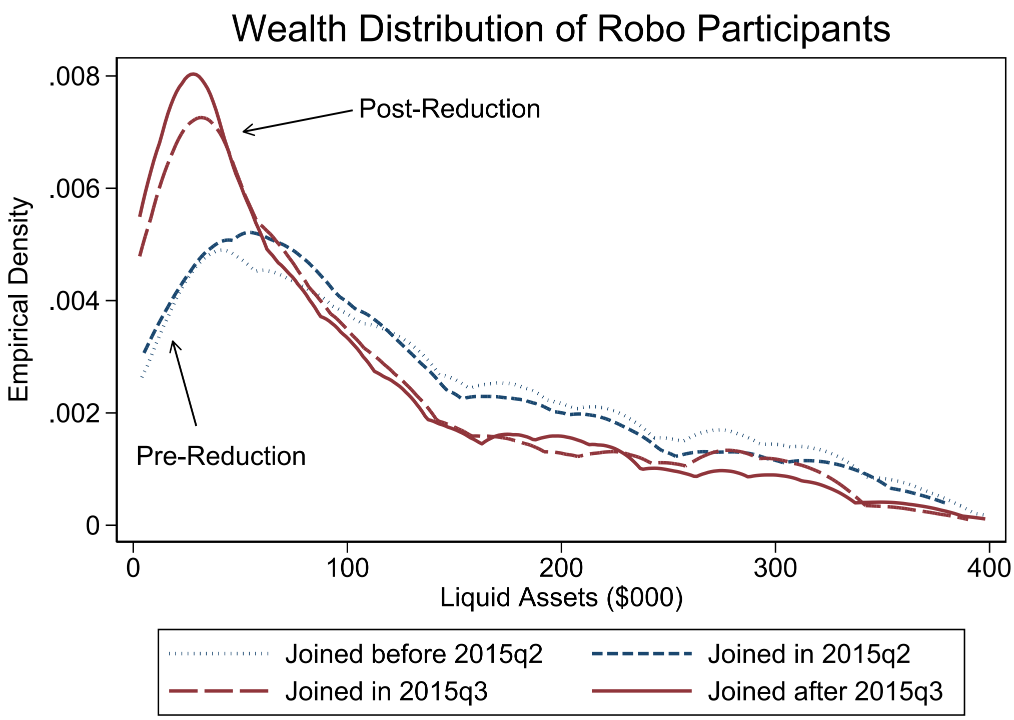

Note: This figure plots the distribution of liquid assets among investors who participated with the robo advisor before the reduction in account minimum (Existing Participants) and who become robo participants after the reduction (New Participants). Liquid assets are defined in Table 1. We remove a very small number of wealthy outliers to fit the distribution on one page. The distribution is calculated using a kernel density.

Figure 3: Change in Representativeness of Robo Wealth Distribution

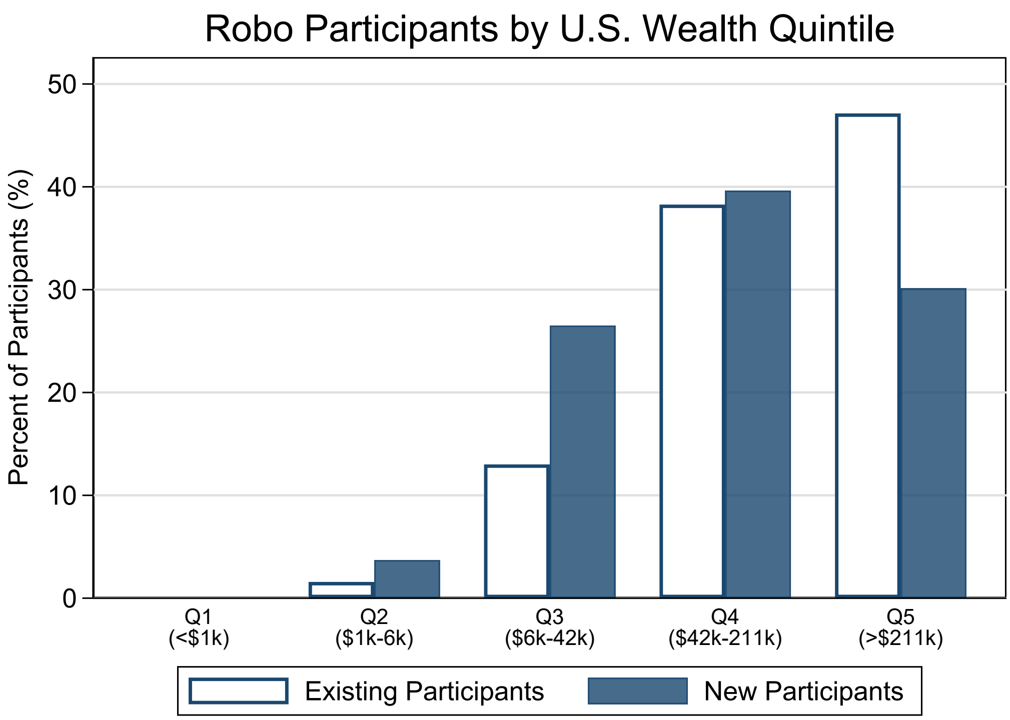

Note: This figure plots the share of robo participants from each quintile of the U.S. wealth distribution. The share is calculated separately for investors who participated before the reduction in account minimum (Existing Participants) and who become participants after the reduction (New Participants).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

Figure 4: Event Study. The Effects of the Reduction in Account Minimum on the Middle Class.

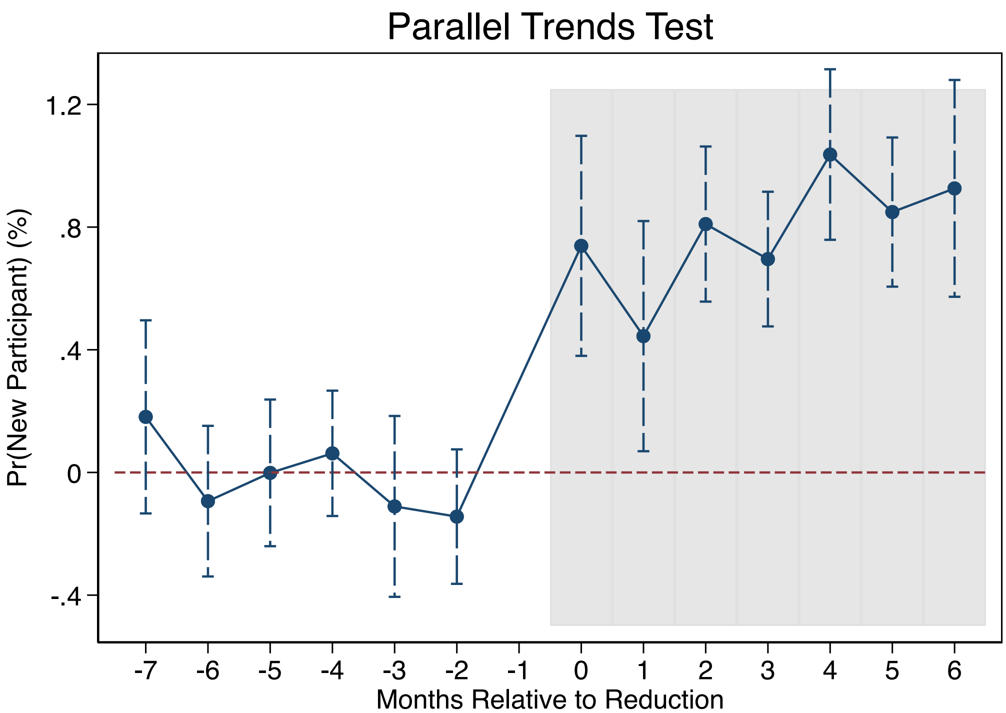

Note: This figure assesses the effect of the reduction in account minimum on middle-class investors' robo participation by estimating the following specification:

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

Subscripts i and t index investor and week. 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 investor 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. Time -1 is the reference month (June 2015). The controls in X_i are the investor's age and an indicator for whether the investor chooses a higher risk tolerance score than that recommended by the robo advisor. The remaining notes are the same as in Table 2.

Figure 5: Event Study. Bunching Behavior by the Middle Class.

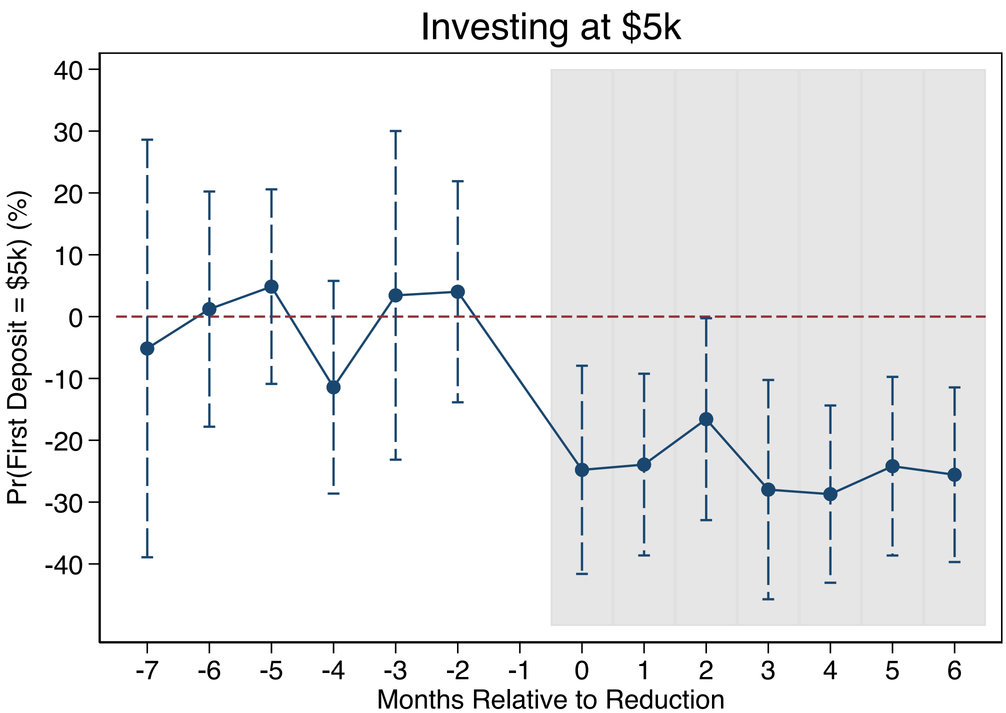

Note: This figure assesses the change in bunching for middle-class investors around the reduction in account minimum by estimating the following specification:

$$\mathbb{1}[\textit{First Deposit}_{i,t} = \$5k] = \beta \textit{Middle}_i + \gamma X_i + \varrho_t + \sum_{m \neq June\,2015} (\mu_m \times \textit{Middle}_i \times \mathbb{1}_{t \in m}) + \sum_{m \neq June\,2015} (\psi_m \times X_i \times \mathbb{1}_{t \in m}) + u_{i,t}.$$

Subscripts i and t index investor and week. The figure shows the estimated coefficients μ_m,which are interpreted as the average difference in the first deposit amount between the middle and the upper classes among all weeks in a given month. Brackets are 95% confidence intervals with standard errors clustered by investor 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. Time -1 is the reference month (June 2015). We do not include investor fixed effects, since the sample includes only the first deposits. The remaining notes are the same as in Table 2.

Figure 6: Advertising Spending by the Robo Advisor

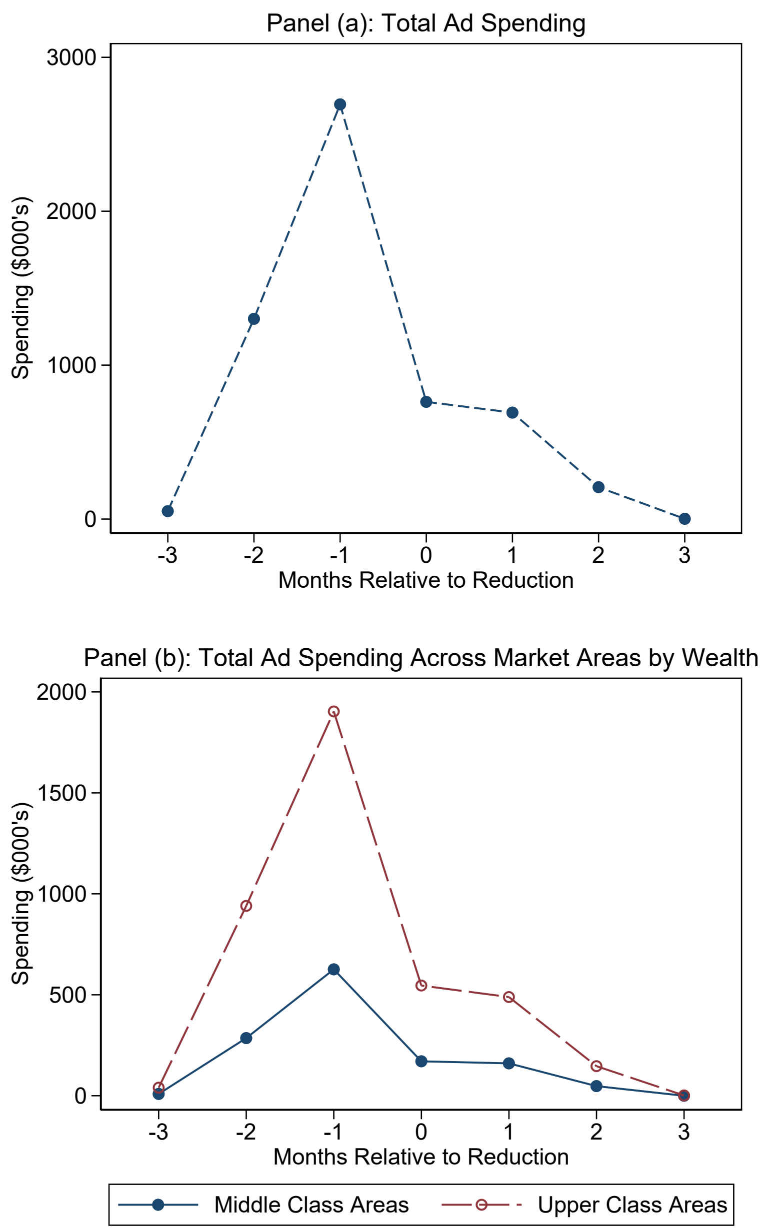

Note: This figure presents the times-series of monthly ad spending by the advisor, surrounding the reduction in account minimum such that time 0 is the month when the reduction was implemented. Panel (a) shows the total ad spending. Panel (b) shows the total spending in middle-class and upper-class Designated Market Areas (DMAs). A DMA is classified as middle-class if its imputed stock market wealth in 2015 is in quintiles two or three of the wealth distribution across all DMAs. DMAs in the top two quintiles are classified as upper-class. Ad spending data are from Vivvix by Kantar. Data on county-level imputed stock market wealth are from Chodorow-Reich, Nenov and Simsek (2021), which we aggregate to the DMA level. Details on data construction are in Appendix A.

Figure 7: Participants across Robo Advisors. Assessing Business Stealing.

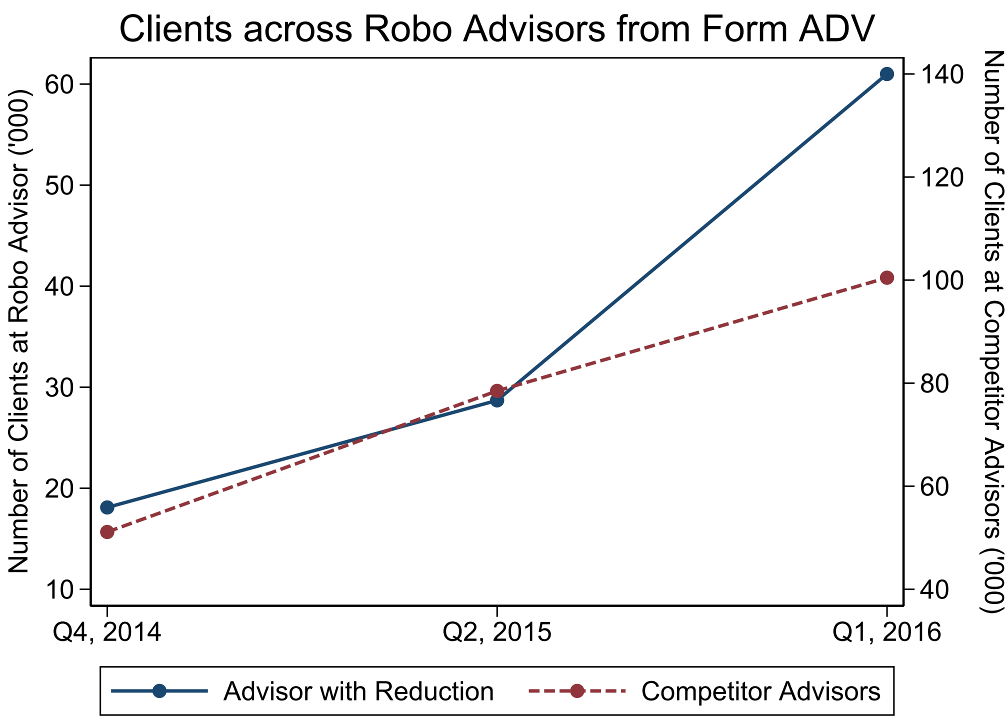

Note: This figure plots the number of clients across robo advisors, in thousands, which assesses whether the reduction increases robo participation or simply reallocates robo participants across advisors. We plot this value separately for the robo advisor that reduced its account minimum, Wealthfront, 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 8: Model Mechanics. Policy Rules under the Previous Minimum of $5,000.

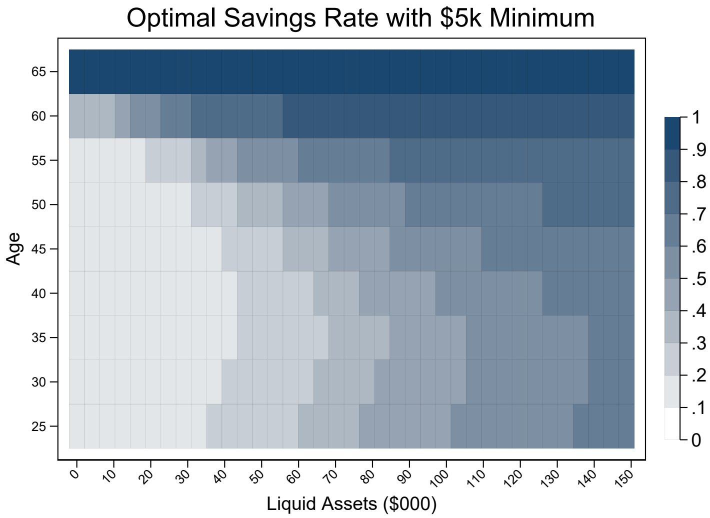

Panel (a)

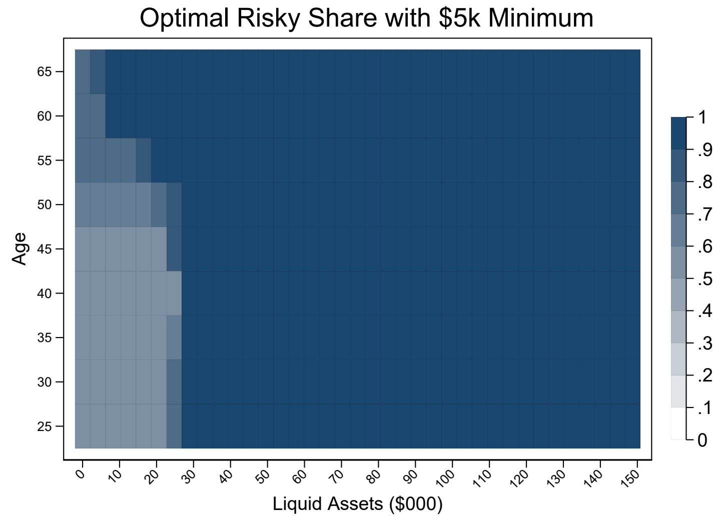

Panel (b)

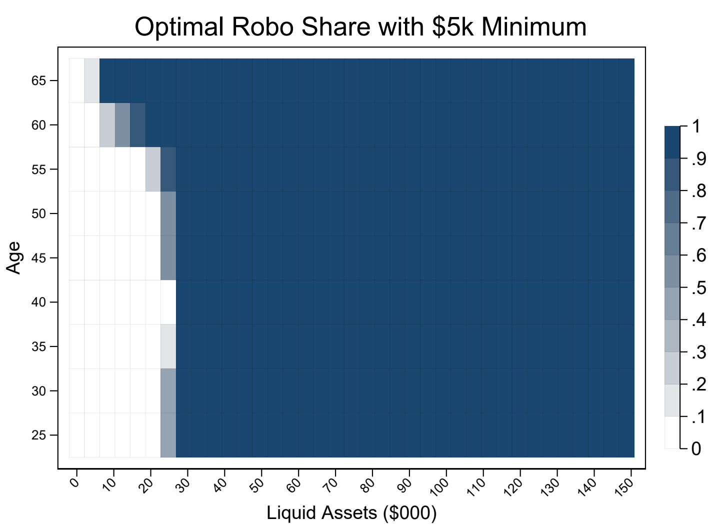

Panel (c)

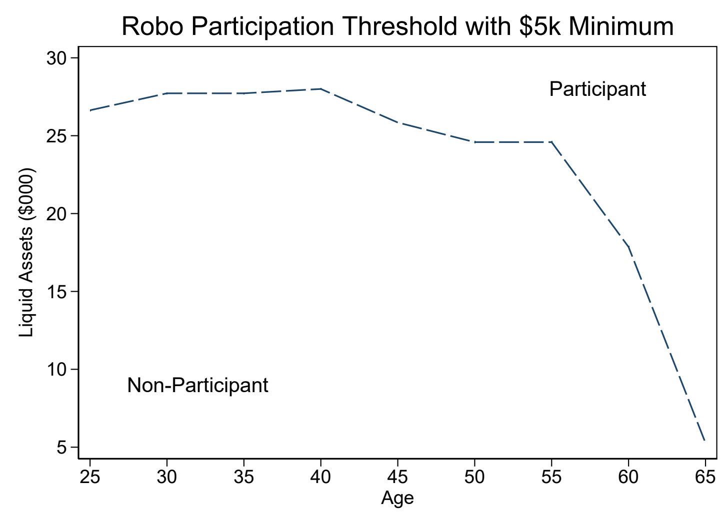

Panel (d)

Note: This figure summarizes investors' optimal participation, portfolio choice, and savings behavior (i.e., policy rules) under a $5,000 account minimum, based on the model in Section 6. Panel (a) plots the savings rate, which is the ratio of the investor's total savings to cash-on-hand. Total savings is the difference between cash-on-hand and consumption. Panel (b) plots the share of the investor's total savings allocated to either the robo portfolio or the self-managed portfolio. Panel (c) plots the share of the investor's savings allocated to the robo portfolio. Panel (d) plots the threshold of liquid assets at which the investor participates with the robo advisor, by 5-year age bin. Policy rules are averaged across years in each age bin. Appendix C contains discussion and details.

Figure 9: Rationalizing the Data with a Benchmark Model of Asset Allocation

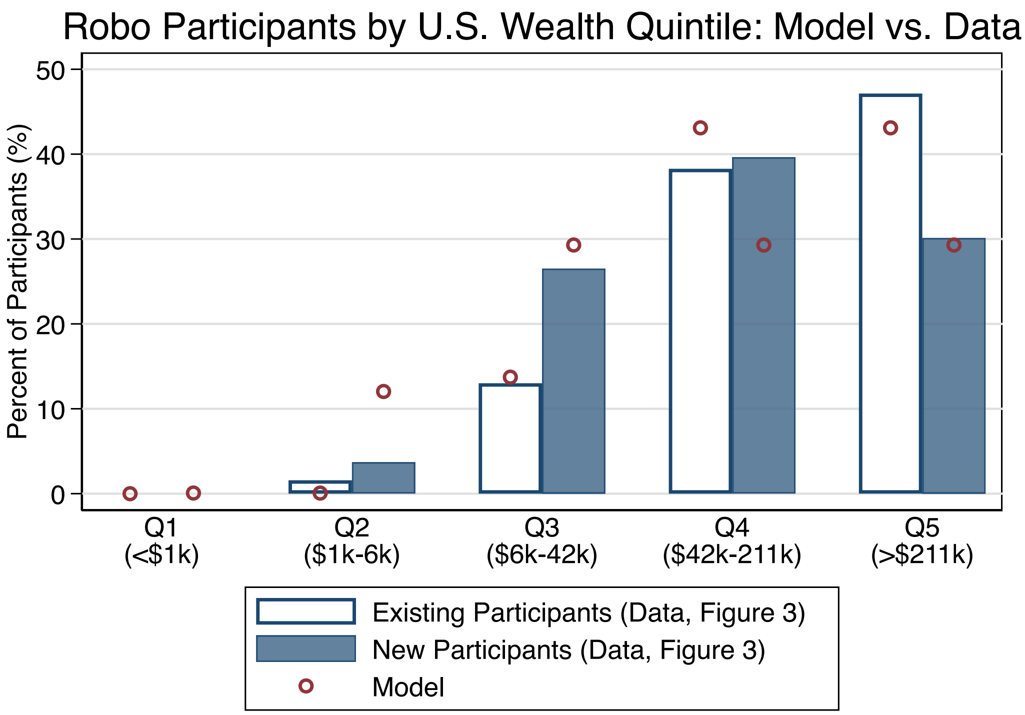

Note: This figure plots the share of robo participants from each quintile of the U.S. wealth distribution, based on the portfolio choice model in Section 6. The shares are shown by the red open circles. Explicitly, figure shows the distribution across wealth quintiles for participants who find it optimal 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 investor-level policy functions across the bins of age and liquid assets that define the state space, weighting by share of the 2016 SCF population within each bin. For reference, the bars plot the empirical share of robo participants from each quintile of the U.S. wealth distribution, already shown in Figure 3. The remaining notes are the same as in Figure 3.

Figure 10: Understanding why Middle-Age Investors Gain More.

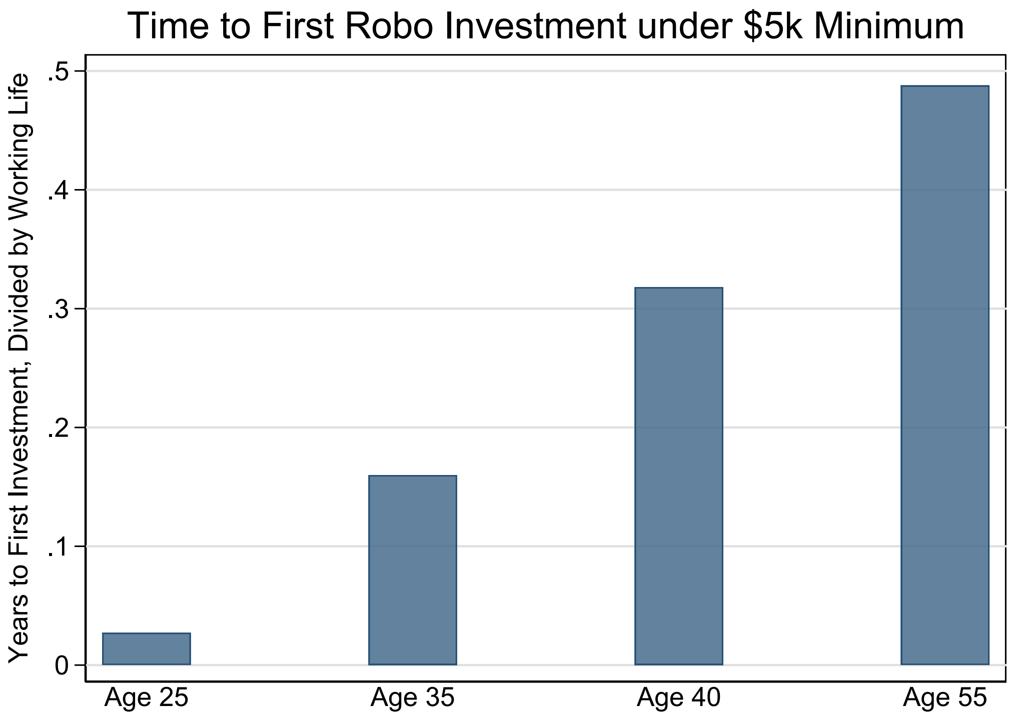

Note: This figure plots the amount of time until a middle-class investor first participates with the robo advisor under an account minimum of $5,000. The amount of time is calculated based on 5,000 simulated paths of the model in Section 6. The outcome variable equals the average number of years until an investor who begins age t with $6,000 in liquid assets becomes a robo participant, where t equals one of the values shown on the x-axis. The number of years until participation is then normalized by the number of years until the investor retires. The starting level of liquid assets of $6,000 equals the 40$^{\text{th}}$percentile of the U.S. wealth distribution. The starting level of permanent income, $\nu_{i,t}$, is chosen to match the average value of permanent income among investors 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 investors starting at age 25 and calculating the distribution of $\nu_{i,t}$ at each subsequent age. Explicitly, $\nu_{i,t} = -0.03 \times (t - 25)$.

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 Investors:								
$Liquid\ Assets_i$ ('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)
$First\ Deposit\ at\ \$5k_i$	0.33	0.47		0.11	0.32		-0.212***	(0.012)
$First\ Deposit < \$5k_i$				0.42	0.49			
$No\ Account\ Closure_i$				0.98	0.15			
$Subsequent\ Deposit_i$				0.71	0.45			
(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)
$First\ Deposit\ at\ \$5k_i$	0.55	0.5		0.09	0.29		-0.460***	(0.028)
$First\ Deposit < \$5k_i$				0.66	0.47			
$No\ Account\ Closure_i$				0.97	0.18			
$Subsequent\ Deposit_i$				0.72	0.45			

Note: This table summarizes investors 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 investor. $Liquid\ Assets_i$ 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 investor income, in thousands of dollars. $Initial\ Deposit_i$ is the value of the investor's initial deposit, in thousands of dollars. Age_i is the investor's age. $High\ Risk\ Tolerance_i$ indicates if the investor 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 investor does not close the account, over the sample period. $Subsequent\ Deposit_i$ indicates if the investor makes a subsequent deposit after her initial one. Most (86%) of subsequent deposits are less than the initial deposit.$First\ Deposit\ at\ \$5k_i$ indicates if the investor's initial deposit equals \$5,000 or is no more than 5% higher. $First\ Deposit < \$5k_i$ indicates if the investor'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 investors who participate with the robo advisor and make a deposit over the period from December 2014 through February 2016. The upper panel summarizes all investors in the sample, and the lower panel summarizes investors 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***	10.560**	14.516***	15.467***		
	(1.357)	(1.339)	(5.116)	(2.966)	(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	2^{nd} or 3^{rd} Quintile	2^{nd} or 3^{rd} Quintile	2^{nd} or 3^{rd} Quintile	2^{nd} or 3^{rd} Quintile	2^{nd} Quintile	Buffer on Middle	2^{nd} or 3^{rd} Quintile	2^{nd} or 3^{rd} 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		
Investor 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	0.076	0.078	0.098	0.008	0.008
Number of Observations	9,349	9,349	9,349	1,901	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 investors into the market for robo advice. Subscripts i and t index investor 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 $Middle_i$. Column (5) uses an indicator for whether i belongs to the second U.S. wealth quintile, after assigning a missing value to investors from the third U.S. wealth quintile. Column (6) uses an indicator for whether i belongs to the second or third U.S. wealth quintile, after assigning a missing value to investors whose liquid assets are within a 10% buffer of the third quintile. 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 investors in the Portfolio Dataset, which includes robo non-participants and a subset of investors in the Deposits Dataset. 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),

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

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 investor fixed effects. Column (8) additionally interacts $Post_t$ with investor controls and a vector of state fixed effects. Wealth consists of liquid assets, defined in Table 1. Wealth quintiles are calculated using the 2016 Survey of Consumer Finances (SCF) dataset. Investor controls X_i are: the log of annual investor income; the investor's age; and an indicator for whether the investor chooses a lower risk tolerance score than that recommended by the robo advisor. Column (4) additionally controls for the log of the outside portfolio's size. Standard errors are clustered by investor (heteroskedasticity robust) in columns (1)-(6), and are two-way clustered by investor and week in columns (7)-(8).

Table 3: Advertising as a Confounding Factor.

$Y_{i,t}$ =	*New Participant*$_{i,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$Middle_i \times Post_t$	0.729***	0.557***	0.611***	0.816***	0.652***	0.689***
	(0.070)	(0.078)	(0.078)	(0.079)	(0.081)	(0.084)
$Middle_i \times Post_t \times Adv_{s,t}$	-0.002	-0.002	-0.005			
	(0.002)	(0.002)	(0.003)			
$Middle_i \times Adv_{s,t}$	-0.000	-0.000	-0.000			
	(0.001)	(0.001)	(0.001)			
$Middle_i \times Post_t \times Adv_{w,s,t}$				-0.028	-0.028	-0.031
				(0.036)	(0.036)	(0.037)
$Middle_i \times Adv_{w,s,t}$				0.012	0.012	0.011
				(0.007)	(0.007)	(0.007)
Investor FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
State × Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × $Post_t$	No	Yes	Yes	No	Yes	Yes
Controls × $Adv_{s,t}$	No	No	Yes	No	No	No
Controls × $Adv_{s,t}$ × $Post_t$	No	No	Yes	No	No	No
$Adv_{w,s,t}$	No	No	No	Yes	Yes	Yes
$Post_t \times Adv_{w,s,t}$	No	No	No	Yes	Yes	Yes
Controls × $Adv_{w,s,t}$	No	No	No	No	No	Yes
Controls × $Adv_{w,s,t}$ × $Post_t$	No	No	No	No	No	Yes
R-squared	0.018	0.018	0.018	0.018	0.018	0.018
Number of Observations	477,005	477,005	477,005	447,050	447,050	447,050

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 investors. Subscripts i and t index investor and week. The baseline regression equation in columns (1)-(3) is of the form

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

where $Adv_{s,t}$ is the robo advisor's ad spending in state s in the month of week t, in thousands of dollars; and$\gamma_{s,t}$ are state-by-week fixed effects. Columns (4)-(6) use a measure of within-state ad spending$Adv_{w,s,t}$: the average spending in state s in the month of week t, separately for middle-class and upper-class Designated Market Area (DMAs), again in thousands of dollars. A DMA is classified as middle-class if its average imputed stock market wealth falls within quintiles two or three of the wealth distribution across all DMAs. DMAs in the top two quintiles are classified as upper-class. The ad spending data are sourced from Vivvix by Kantar. Details on data construction are outlined in Appendix A and Section 5.2. The remaining notes are the same as in Table 2.

Table 4: 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***
Expected Return	0.080	0.102	0.023***	0.078	0.101	0.023***
Total Volatility	0.209	0.137	-0.071***	0.196	0.134	-0.062***
Idiosyncratic Volatility	0.146	0.034	-0.111***	0.138	0.033	-0.104***
Robo-Sharpe Expected Return	0.154			0.147		
Return Loss	-0.074			-0.069		

Note: This table summarizes portfolios that investors 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; *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 two bond factors: the monthly percent changes in Barclays' U.S. and global bond indices (Fama-French with Bond). Expected returns for a portfolio are calculated as the dot product of the portfolio's loading, shown in equation (9), and the risk prices of the factors, reported in Appendix Table A14. Columns (1)-(2) report the mean across portfolios for investors 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 investors who consult the robo advisor for a free portfolio review. Of these investors, 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 5: Model Parameters

Parameter	Value	Source
(a) Portfolio Parameters:		
Idiosyncratic Volatility ($\sigma_{\epsilon,i}$)	Summarized in Table 4	Portfolio Dataset
Factor Loadings (β_i)	Summarized in Appendix Table A13	Portfolio Dataset
(b) Asset Pricing Factors:		
Market Factor, Mean	0.076	CRSP
Market Factor, Volatility	0.147	CRSP
Fama-French Factors	Appendix Table A14	French
Bond Factors	Appendix Table A14	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 (σ_ξ)	0.271	CGM
Correlation with Financial Return (ρ^Y)	0.003	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]	Standard
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 correlation between log labor income and financial returns of $\rho^Y = 0.3\%$ corresponds to a loading of 0.007; the probability of labor income disaster is calculated as the share of investors 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; investors 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 an investor 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; 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 4. Appendix A has details on these data sources.

Table 6: Rationalizing the Data with a Benchmark Model of Asset Allocation

	Model	Data	Source
	(1)	(2)	(3)
(a) Baseline Model, Robo Participation:			
Pre-Reduction Share from Middle-Class	14.9%	12.6%	Deposits Dataset
Post-Reduction Share from Middle-Class	43.1%	28.4%	
(b) Baseline Model, Risky Share:			
Risky Share, Existing Middle-Class	32.3%	31.1%	Deposits Dataset
Risky 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 previous minimum and the reduced minimum; in the data, we compute the ratio of robo investment to liquid assets among middle class investors 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 investors who participate in the stock market, calculated as follows: in the model, we compute the share of middle-class investors 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 investors 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 investors incur a per-period cost of $\kappa = \$200$ when investing in the self-managed portfolio (Appendix 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 investor-level policy functions across the bins of age and liquid assets that define the state space, weighting by the share of the 2016 SCF within each bin. The remaining notes are the same as in Table 2.

Table 7: Distribution of Welfare Gains and Channels that Drive the Gains

	Increase in Lifetime Consumption			
		Middle Class		
	Pooled	Age 25-35	Age 36-55	Age 56-65
	(1)	(2)	(3)	(4)
(a) Welfare Gain From the Reduction:	0.77%	0.58%	0.59%	1.68%
(b) Gain from Reduction on Counterfactual Robo Portfolios:				
(i) Asset Allocation Counterfactuals				
Same Idiosyncratic Risk as Self-Managed Portfolio	0.22%	0.11%	0.10%	0.80%
Same Market Loading as Self-Managed Portfolio	0.71%	0.54%	0.51%	1.60%
Same Value and Size Loadings as Self-Managed Portfolio	0.46%	0.32%	0.26%	1.31%
Same Bond Loadings as Self-Managed Portfolio	0.20%	0.15%	0.10%	0.58%
(ii) Glide Path Counterfactuals				
No Adjustment for Age	0.34%	0.34%	0.20%	0.73%
No Adjustment for Wealth	0.37%	0.33%	0.23%	0.88%
No Adjustment for Age or Wealth	0.24%	0.26%	0.15%	0.44%
(c) Gain Compared to Hypothetical Self-Managed Portfolios:				
Compared to Stock Index Fund	0.72%	0.47%	0.52%	1.84%
Compared to Target Date Fund	0.23%	0.16%	0.17%	0.47%

Note: This table summarizes the average welfare gain for investors 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 an investor's expected lifetime utility by the same amount as the reduction, as in equation (10). 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 investor i to equal those of the self-managed portfolio for investor i. The constraints are: the same idiosyncratic 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 parameters of the robo portfolio ($\beta^{\mathcal{A}}_{i}, \sigma^{\mathcal{A}}_{\epsilon,i}$) by investor 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). Panel (c) performs a similar calculation as in panel (a) after specifying a hypothetical self-managed portfolio as the comparison group: Vanguard's Total Stock Market ETF (VTI); and the Vanguard TDF in which an investor of a given age t in 2015 and retirement age 65 would invest. The remaining notes are the same as in Table 2.

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. Thus, we observe the same information as would an analyst working for the advisor. 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.[32]

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 investor's annual income, state of residence, liquid assets, recommended and selected risk tolerance score, and investor 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 investor'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 $\textit{Risk Averse}_i$ in the text. Only 3% of investors 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 0.15.

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

[32] The reduction technically occurs in the second week of July, and so, to be as precise as possible, we classify the first week of July as if it were still June in our regressions. This technicality has no bearing on the results.

'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 investors in the deposits dataset as well as a set of robo non-participants. It contains snapshots of investors' portfolio holdings in an outside, traditional brokerage account. This information is paired with the portfolio holdings of the investor'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 investors' non-robo accounts. While the advice algorithm ran, our data provider would ask the investor 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 investor 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 investor 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 4.

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. retail investors. 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 A4 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 Advertising Datasets

We use advertising data from Kantar, specifically Kantar's Vivvix dataset. Kantar tracks traditional media advertisements (e.g., television, print, radio) and, increasingly, digital advertisements. Kantar measures traditional advertisements by tracking which advertisements are conveyed in cable or broadcast television or radio signals. Gordon et al. (2020) discuss the relationship between a firm's traditional and digital advertisements, noting that they tend to be highly correlated.

The unit of observation in the Kantar dataset is a tuple defined by: company; time of advertising; method of advertising (e.g., television, newspaper); geographic location, if relevant for the given method; and product being advertised. We observe total dollar spending for each tuple. The only product observed for the robo advisor is their baseline service, described as "Wealthfront Online".

Kantar measures geography at the level of the Designated Market Area (DMA). Loosely, a DMA corresponds to a CBSA. We map each DMA to a set of counties using the crosswalk file from Gentzkow (2006) and Gentzkow and Shapiro (2008). Precisely, this crosswalk file maps county FIPS codes to the DMA name in a related advertising dataset, Nielsen's Ad Intel dataset. The DMA names in Nielsen are effectively the same as in Kantar after small adjustments in spelling. This mapping enables us to merge our advertising data with data on county-level imputed wealth from Chodorow-Reich, Nenov and Simsek (2021). We aggregate imputed wealth across counties by taking the average for all counties associated with a DMA, weighting by population.

We cross-reference some of the statistics from the Kantar dataset with Nielsen's Ad Intel dataset, as noted in footnote 19. The Nielsen dataset has the same format as the Kantar dataset, except that we do not observe the DMA associated with the advisor's television ad spending over our period of analysis.

A.3.4 Imputed Stock Market Wealth by County

Our data on imputed county-level stock market wealth come from Chodorow-Reich, Nenov and Simsek (2021), who apply a county-specific capitalization rate to total dividend income in a county from the IRS SOI Tax Stats county-level dataset. The baseline capitalization rate is the aggregate price-dividend ratio from the CRSP dataset. Then, Chodorow-Reich, Nenov and Simsek (2021) adjust this baseline capitalization rate for each county according to: the county's average age; and the share of wealth that is taxable. The age adjustment involves calculating the weighted average price-dividend ratio for bins defined age, using data on the positions of retail investors from Barber and Odean (2000). The weights used to calculate this average come from the joint distribution of wealth by age from the SCF dataset. Each county is assigned an age-specific capitalization rate. The adjustment for nontaxable income involves projecting total stock market wealth on taxable stock market wealth and household demographic characteristics using the SCF dataset. This projection is used to map taxable stock market wealth, obtained from the age adjustment just described, onto total stock market wealth. Finally, Chodorow-Reich, Nenov and Simsek (2021) perform this imputation on a county-by-year panel, and we simply use the imputed stock market wealth from 2015. We obtain the data from the replication file for Chodorow-Reich, Nenov and Simsek (2021).

A.3.5 Other Auxiliary Datasets

We use the Center for Disease Control's mortality tables to calibrate the survival probabilities 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 investor 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 investor i's liquid assets fall within the second or third U.S. quintile of liquid assets. Investor 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 investor 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 investors and equals 0 for investors who participated before July 7, 2015.

- *Initial Deposit*$_i$: This variable is the initial deposit with the robo advisor made by investor i, based on the deposits dataset.

- *Income*$_i$: This variable is annual income for investor i, based on the deposits dataset.

- *Age*$_i$: This variable is the age of investor i, based on the deposits dataset.

- *Risk Averse*$_i$: This variable indicates if investor 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 investor 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 investor i's initial deposit with the robo advisor is between $5,000 and $5,250, based on the deposits dataset.

B Econometric Appendix

First, we provide the details for aggregation exercise referenced in Section 4.3. Next, we explain how the estimation of equation (2) can be affected by measurement error (B.2). Lastly, we provide a formal argument for how the advertising robustness in Table 3 can accommodate targeted advertising by wealth.

B.1 Details on Aggregation Exercise to Interpret the Magnitude

We provide details on the aggregation exercise referenced in Section 4.3, 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 investors 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 investors 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 investors do not experience a relaxation of constraints and, thus, $\mu = 0$. Under this counterfactual, the overall number of robo participants grows at the rate

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

or, equivalently,

$$g^C = \frac{\mathbb{E}\left[\textit{New Participant}_i\right] - \mu\mathbb{E}\left[\textit{Middle}_i\right]}{1 - \left(\mathbb{E}\left[\textit{New Participant}_i\right] - \mu\mathbb{E}\left[\textit{Middle}_i\right]\right)}. \tag{B5}$$

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

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

which we can rewrite as

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

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 A10 summarizes various calculations of η and of the analogous statistic for growth in middle-class investors' 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.[33] 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 4.3.

B.2 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

$$\widehat{\mu} = \mu\left(1 - \frac{\operatorname{Var}\left[\varepsilon_i\right] + \mathbb{E}\left[Middle_i \times \epsilon_i\right]}{\operatorname{Var}\left[\widehat{Middle}_i\right]}\right) + \frac{\mathbb{E}\left[u_i \times \varepsilon_i\right]}{\operatorname{Var}\left[\widehat{Middle}_i\right]}. \tag{B8}$$

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.3 Interpreting the First-Difference Estimator

Section 4 wrote that one can interpret the coefficient in the first difference estimator, μ, as the share of robo participants whom the reduction has caused to participate. We now substantiate this remark. Note that the reduction causes an investor to participate if she finds it optimal to become a robo investor under the \$500 minimum but not under the \$5,000 minimum. Let C_i indicate such an investor. The share of eventual participants whom the reduction causes to participate is thus: $\Pr[C_i = 1 \mid Participant_{i,1} = 1]$. Under the assumption that the reduction does not affect existing middle-class participants or upper-class participants:

$$\Pr[C_i = 1 \mid Part_{i,1} = 1] = \Pr[New_i = 1 \mid Part_{i,1} = 1] - \Pr[New_i = 1, Middle_i = 0 \mid Part_{i,1} = 1] = \mu, \tag{B9}$$

[33] 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.

where $Part_{i,t}$ abbreviates$Participant_{i,t}$. The last equality follows from equation (2) after omitting controls X_i for brevity.

C Model Appendix

We describe the solution of the model in Section 6, derive the welfare measure studied in Section 7.2, provide details on the model extensions also studied in Section 7.2, and structurally estimate the model's preference parameters as referenced in Section 6.2.

C.1 Model Solution

We first describe the labor income process. Then, we state the problem. Finally, we describe the numerical solution algorithm.

C.1.1 Details on Labor Income Process

Investors 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

$$\log(Y_{i,t}) = f_i + \xi_{i,t} + \nu_{i,t} + \Xi_t, \tag{C1}$$

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

$$\nu_{i,t} = \nu_{i,t-1} + \upsilon_{i,t}, \tag{C2}$$

where $\upsilon_{i,t}$ is normally distributed with mean zero and volatility of σ_υ. The aggregate component, Ξ_t, covaries with financial returns according to

$$\Xi_t = \beta^Y \left(\bar{R}_t^A - \mathbb{E}[\bar{R}_t^A] \right) \tag{C3}$$

where $\bar{R}_t^A$ is the robo portfolio return at time t, averaged across investors of the same age; $\mathbb{E}[\bar{R}_t^A]$ is the expected value of this return, evaluated over the distribution of F_t; and β^Y is the loading of the robo portfolio's abnormal return in year t. Cocco, Gomes and Maenhout (2005) use the Panel Study of Income Dynamics and the method of Carroll and Samwick (1997) to estimate $\sigma_\nu = 0.103$ and $\sigma_\xi = 0.271$. They estimate a correlation coefficient between log income and the stock market's abnormal return of $\rho^Y = 0.6\%$ for their baseline income group, which maps approximately to $\beta^Y = 0.015$, but they use a correlation coefficient of zero in their baseline analysis. Accordingly, we take the average, corresponding to $\rho^Y = 0.3\%$ and $\beta^Y = 0.007$ as shown in Table 5.

Several studies find that including income skewness improves the performance of life cycle models (e.g., Guvenen, Ozkan and Song (2014); Bagliano, Fugazza and Nicodano (2018); Catherine (2022)). Following Carroll (1997) and Cocco, Gomes and Maenhout (2005), we incorporate skewness by introducing a disaster state in which investors receive zero labor income for one year. Such disasters occur with probability ϕ. In years without a disaster, labor income is given by equation (C1).

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, investors 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. We state the retirement problem shortly, in Appendix C.1.2.

C.1.2 Consolidated Problem

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

$$C_{i,t} = [1 - \alpha^f_{i,t} - \alpha^{\mathcal{S}}_{i,t} - \alpha^{\mathcal{A}}_{i,t}]W_{i,t}. \tag{C4}$$

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

$$\alpha^f_{i,t} \geq 0, \ \ \alpha^{\mathcal{S}}_{i,t} \geq 0, \ \ \alpha^{\mathcal{A}}_{i,t} \geq 0, \tag{C5}$$

$$1 - \alpha^f_{i,t} - \alpha^{\mathcal{S}}_{i,t} - \alpha^{\mathcal{A}}_{i,t} \geq 0, \tag{C6}$$

$$\alpha^{\mathcal{A}}_{i,t} = 0 \ \text{ or } \ \alpha^{\mathcal{A}}_{i,t} \geq \frac{M}{W_{i,t}}. \tag{C7}$$

Constraint (C5) rules out borrowing and shorting, which we subsequently relax in Appendix C.4. Constraint (C6) ensures nonnegative consumption. Both of these constraints are standard. The third constraint, (C7), 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 ($W_{i,t}$). As noted by Cocco, Gomes and Maenhout (2005), the problem is homogeneous in permanent labor income, $v_{i,t}$, allowing us to remove it from the set of state variables. Liquid assets evolve according to

$$W_{i,t+1} = \left[\alpha^f_{i,t}(1 + R^f) + \alpha^{\mathcal{S}}_{i,t}(1 + R^{\mathcal{S}}_{i,t+1}) + \alpha^{\mathcal{A}}_{i,t}(1 + R^{\mathcal{A}}_{i,t+1})\right] W_{i,t} + Y_{i,t+1}. \tag{C8}$$

As in the text,

$$R^{\mathcal{P}}_{i,t} = \beta^{\mathcal{P}}_i F_t + \epsilon^{\mathcal{P}}_{i,t}, \tag{C9}$$

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

$$V_t(W_{i,t}) = \max_{\alpha^f_{i,t}, \alpha^{\mathcal{S}}_{i,t}, \alpha^{\mathcal{A}}_{i,t}} \left\{ \frac{{C_{i,t}}^{1-\gamma}}{1-\gamma} + \delta p_t \mathbb{E}_t \left[V_{t+1}(W_{i,t+1}) \right] \right\} \tag{C10}$$
$$s.t. (C4)\text{-}(C9)$$

for $t_0 \leq t \leq \underline{T}$. For $t > \underline{T}$, investors solve

$$V_{\underline{T}+1}(W_{i,\underline{T}+1}) = \max_{\left\{\alpha^f_\tau\right\}} \sum_{\tau=\underline{T}+1}^{\tau=\overline{T}} (\delta p_{\underline{T}+1})^{\tau-\underline{T}-1} \frac{{C_{i,\tau}}^{1-\gamma}}{1-\gamma} \tag{C11}$$
$$s.t.$$
$$0 \leq \alpha^f_\tau \leq 1 \tag{C12}$$
$$W_{i,\tau+1} = \alpha^f_{i,\tau}(1 + R^f)W_{i,\tau} \tag{C13}$$
$$C_{i,\tau} = [1 - \alpha^f_{i,\tau}]W_{i,\tau}. \tag{C14}$$

Indirect utility has the familiar form

$$V_{\underline{T}+1}(W_{i,\underline{T}+1}) = B\frac{{W_{i,\underline{T}+1}}^{1-\gamma}}{1-\gamma}, \tag{C15}$$

with

$$B = \sum_{\tau=\underline{T}+1}^{\overline{T}} \delta^{\tau-\underline{T}-1} \left[\frac{1-\chi}{1-\chi^{\overline{T}-\underline{T}}}\right] \left[\delta(1+R^f)\right]^{\frac{\tau-\underline{T}-1}{\gamma}}, \tag{C16}$$

$$\chi = \delta^{\frac{1}{\gamma}}(1+R^f)^{\frac{1-\gamma}{\gamma}} \tag{C17}$$

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

C.1.3 Numerical Algorithm

Our numerical algorithm is standard and follows the methods typically used in benchmark portfolio choice models. First, we solve equation (C10) 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. The simulation used to produce Figure 10 uses a more fine grid with increments of 0.01 in liquid assets on a log scale. This more-fine grid is also used for Figure 8. We 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: $F_t, \epsilon^{\mathcal{S}}_{i,t}, \epsilon^{\mathcal{A}}_{i,t}, \overline{\nu}_{i,t}$, with $\overline{\nu}_{i,t} \equiv \xi_{i,t} + \nu_{i,t}$. 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 (C5)-(C8). As mentioned in the text, we simplify the model's computational complexity by assuming investors 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 investors must maintain a minimum balance of M with the automated asset manager, whereas, in reality, investors 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 an investor's robo investment. Namely, under an initial deposit requirement, investors 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 7. Intuitively, the high expected return on the robo portfolio makes cases of a top-up quite 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 $\left\{\alpha^f_\tau(W_{i,\tau}), \alpha^S_\tau(W_{i,\tau}), \alpha^{\mathcal{A}}_\tau(W_{i,\tau})\right\}$ that we use in the positive and welfare analyses of Section 7.

C.2 Discussion of Policy Functions and Life Cycle Intuition

Figure 8 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 portfolio 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 panel (d). We plot the minimum threshold of liquid assets at which the investor decides to participate with the robo advisor as a function of investor age. This threshold remains relatively stable at around \$27,000 until the investor 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 investors 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 investors are willing to invest a larger share of their liquid assets with the robo advisor. This reflects how investors' 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 panel (a) of Figure 8.

In benchmark models, investors 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 investors 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 investors age. For example, recall how robo portfolios for older or less-wealth investors have a stronger tilt towards bonds and away from stocks (Figure 1). This glide path feature incentivizes older investors to invest a larger share of their wealth in the robo portfolio. In fact, panel (c) of Figure 8 implies that, beginning around age 40, the optimal robo share increases as investors age, holding liquid assets fixed. This increase in robo share drives an increase in risky share, as we now describe.

Panel (b) of Figure 8 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 investor 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 our baseline setup does not feature fixed costs of stock market participation, the optimal risky share is always positive, even for very low levels of wealth. Note, however, that most investors in the middle class select a risky share well below 100%. Benchmark models, by contrast, predict a 100% risky share for almost all investors. This difference arises because investors in benchmark models invest in the total stock market, and, thus, have a relatively low level of idiosyncratic risk. By contrast, less-wealthy investors in our setting must invest in the self-managed portfolio, which, per Table 4, has a relatively low Sharpe ratio. This underdiversification reduces their demand for risky assets. However, once liquid assets reaches around \$27,000, investors begin to invest in the robo portfolio, and the risky share jumps to 100%.

Lastly, we point out three features of the savings rate policy function in panel (a) that accord with benchmark life cycle models. First, as already described, the savings rate begins to increase in age after the investor 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 investor wishes to accumulate a buffer against idiosyncratic labor income shocks. Such a motive 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 investor's liquid assets reach some threshold. Benchmark models feature no such jump. Moreover, a personalized risk profile by both age and wealth increases older investors' and less-wealthy investors' demand for risky assets beyond what a benchmark model would predict.

C.3 Welfare Measure

Repeating from the text, we measure investor 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 investor 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

$$\mathbb{E}_t\left[\sum_{\tau=t}^{\tau=\overline{T}}\delta^{\tau-t}\left(\prod_{j=t}^{j=\tau-1}p_j\right)\frac{\left((1+q_i)\underline{C}_{i,\tau}\right)^{1-\gamma}}{1-\gamma}\right]=\mathbb{E}_t\left[\sum_{\tau=t}^{\tau=\overline{T}}\delta^{\tau-t}\left(\prod_{j=t}^{j=\tau-1}p_j\right)\frac{\left(\overline{C}_{i,\tau}\right)^{1-\gamma}}{1-\gamma}\right]\equiv\overline{V}_i, \tag{C18}$$

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

$$\underline{V}_i\equiv\mathbb{E}_t\left[\sum_{\tau=t}^{\tau=\overline{T}}\delta^{\tau-t}\left(\prod_{j=t}^{j=\tau-1}p_j\right)\frac{\left(\underline{C}_{i,\tau}\right)^{1-\gamma}}{1-\gamma}\right] \tag{C19}$$

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

$$q_i=\left(\frac{\overline{V}_i}{\underline{V}_i}\right)^{\frac{1}{1-\gamma}}-1, \tag{C20}$$

as shown in equation (10). Note that equation (C20) 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 Extensions and Robustness

We describe the model extensions in Table A18.

C.4.1 Participation Costs

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, investors incur a per-period cost of κ = \$200 when holding the self-managed portfolio, which we obtain from the structural estimation described in Appendix C.5. Explicitly, flow utility becomes: $\frac{(C_{i,t}-\kappa \mathbf{1}\cdot[\alpha^S_{i,t}>0])^{1-\gamma}}{1-\gamma}$.

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) of Table A18. In relation to Table 7, 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 investors 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 $\kappa = \$700$, also shown in Table A18.

C.4.2 Borrowing

Next, we allow investors to borrow up to $b = 30\%$ of liquid wealth. Explicitly, constraint (C5) 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 investors borrow equals 12%, the average interest rate on credit card debt in 2015. The resulting welfare gain in panel (a) of Table A18 equals 0.8%, almost equivalent to the baseline gain. This similarity partly reflects a precautionary savings motive, which discourages middle-class investors from borrowing. Moreover, borrowing to invest in the robo portfolio would be suboptimal given the high interest rate on credit card debt.

C.4.3 Calibration of Portfolio Parameters

We recalculate welfare implications under a different calibration of portfolio parameters. We do so by using the relationship between the portfolio characteristics of Swedish investors 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) of Table A18 shows that the welfare gains under these alternative calibrations are comparable to the baseline.

C.4.4 Calibration of Preference Parameters

We reproduce the main results under the following parameterizations: structurally estimated preference parameter values, based on the estimation described shortly in Appendix C.5; and a high discount factor of $\delta = 0.85$. Panel (c) of Table A18 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 ($\gamma = 5.1$).

The last row of panel (c) recalculates the welfare gain when the correlation between log labor income and the robo portfolio's abnormal return equals 80%, corresponding to $\beta^Y = 3.3$.This extension results in only a very small welfare loss, which may seem counterintuitive. Indeed, the strong positive correlation raises the total risk of investing with the robo advisor, which would tend to reduce the welfare gain. However, by the same logic, this correlation reduces an investor's unconstrained-optimal risky share, which makes investing under a lower minimum especially valuable. On net, the two forces approximately cancel out, leading to almost the same results as in our baseline calculation.

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.[34] 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

$$(\hat{\gamma}, \hat{\delta}, \hat{\kappa}) = \arg \min_{(\gamma,\delta,\kappa)} (\tilde{\Theta}(\gamma, \delta, \kappa) - \hat{\Theta})' \Psi (\tilde{\Theta}(\gamma, \delta, \kappa) - \hat{\Theta}), \tag{C21}$$

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 A19.

[34] 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 within the window from July 1990 through February 2016. Following Calvet et al. (2007), we restrict the set of securities to stocks and funds traded on an exchange. We obtain data on the returns for each security k from the Center for Research in Security Prices (CRSP). Once we have the estimated the loadings $\hat{\beta}_k^F$ from equation (D1) for model F, it is straightforward to compute the idiosyncratic volatility and factor loadings for investor i's self-managed and robo portfolios, as in equation (9).

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, net of the risk-free rate; R_m^{HML} is the spread in monthly return between high book-to-market stocks and low book-to-market stocks (HML); and R_m^{SMB} denotes the spread in monthly return between stocks with a small market capitalization and a big market capitalization (SMB). We measure R_m^M, R_m^{HML}, and R_m^{SMB} using data from Ken French's website, specifically the "Fama/French 3 Factors for Developed Markets" series. The two bond factors are the return on U.S. and global bonds, denoted R_m^{USB} and R_m^{GLB}. We measure R_m^{USB} and R_m^{GLB} using the monthly percent changes in the Bloomberg-Barclays Aggregate U.S. Bond Index Unhedged (R_m^{USB}) and in the Bloomberg-Barclays Global Aggregate Bond Index Unhedged (R_m^{GLB}). Data on these indices come from the robo advisor.[35]

We calibrate the factor covariance matrix and risk prices as follows. The covariance matrix is that of: (i) the return on the CRSP Value-Weighted Index minus the one-month Treasury yield; (ii) the HML return for the U.S. market from Ken French's website; (iii) the SMB return for the U.S. market from Ken French's website; (iv) the return on the Bloomberg-Barclays Aggregate U.S. Bond Index Unhedged; and (v) the return on the Bloomberg-Barclays Global Aggregate Bond Index Unhedged. The covariance matrix is calculated using the subperiod from January 1990 through February 2016 because we only observe the bond factors over that subperiod. We calculate risk prices using the sample means of the series (i)-(v),

[35] The bond indices used to estimate equation (D4) differ slightly from the actual Bloomberg-Barclays indices because of an artefact in the delivery of these indices in January 2017. We later cross-referenced the actual indices with the indices that we work with. The monthly percent change in our indices is on average 0.002 lower than with the actual indices, and the standard deviation of this difference is 0.002. We interpret this noise as a manifestation of the Roll (1977) critique that applies to factor models more generally. Econometrically, this is another argument in favor of using a model to estimate expected returns, since such noise biases the estimated mean of the factor but not necessarily the estimated loading of a security on that factor.

taken over January 1960 through December 2017 for (i)-(iii) and over January 1990 through February 2016 for (iv)-(v). The corresponding covariances and risk prices are shown in Appendix Table A14.

Additional Figures and Tables

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

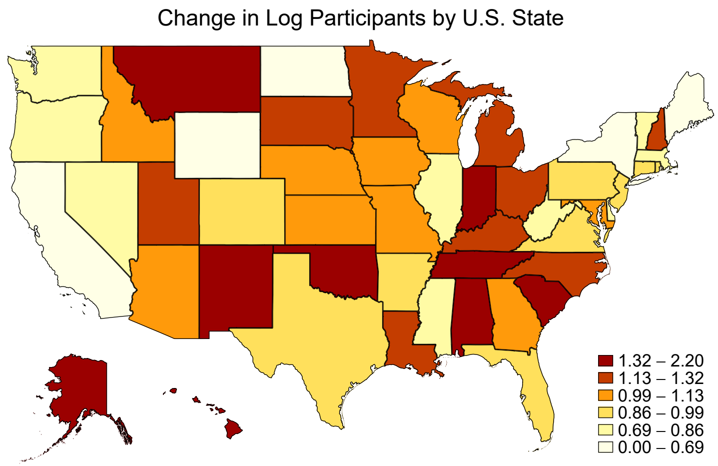

Note: This figure plots the change in the log of the number of robo participants from each U.S. state. The change is from the pre-reduction period (December 1, 2014 to July 7, 2015) to the post-reduction period (July 7, 2015 to February 29, 2016).

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

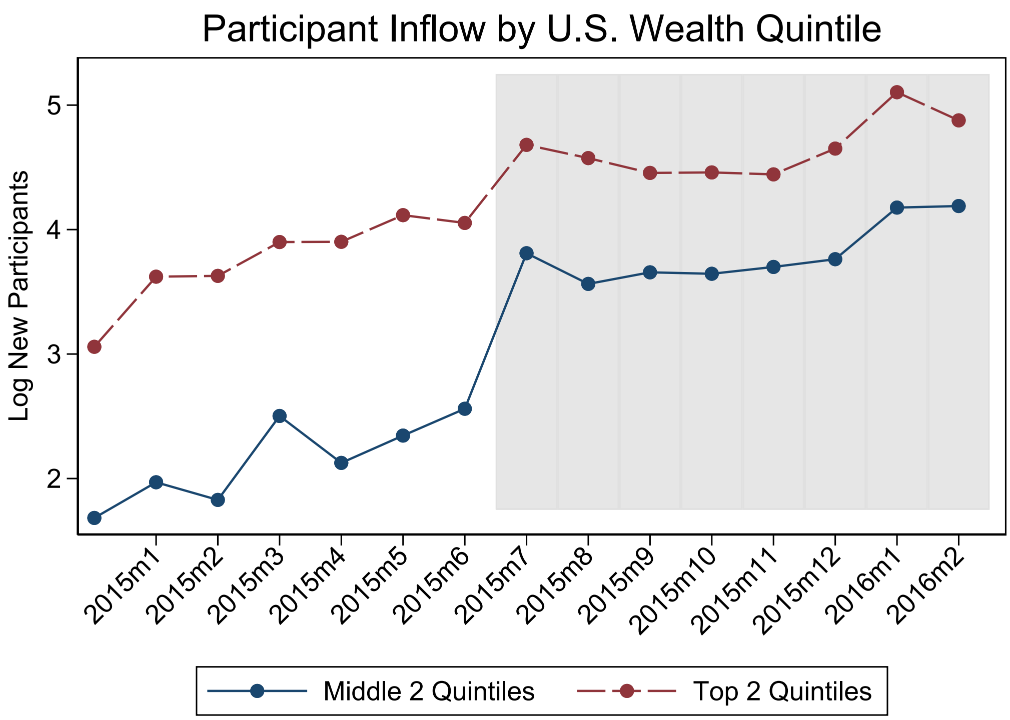

Note: This figure plots the log number of new robo participants from the second and third U.S. wealth quintiles and the fourth and fifth U.S. wealth quintiles, by month. The remaining notes are the same as in Figure 3.

Figure A3: Robustness of the Change in Representativeness to Portfolio Dataset

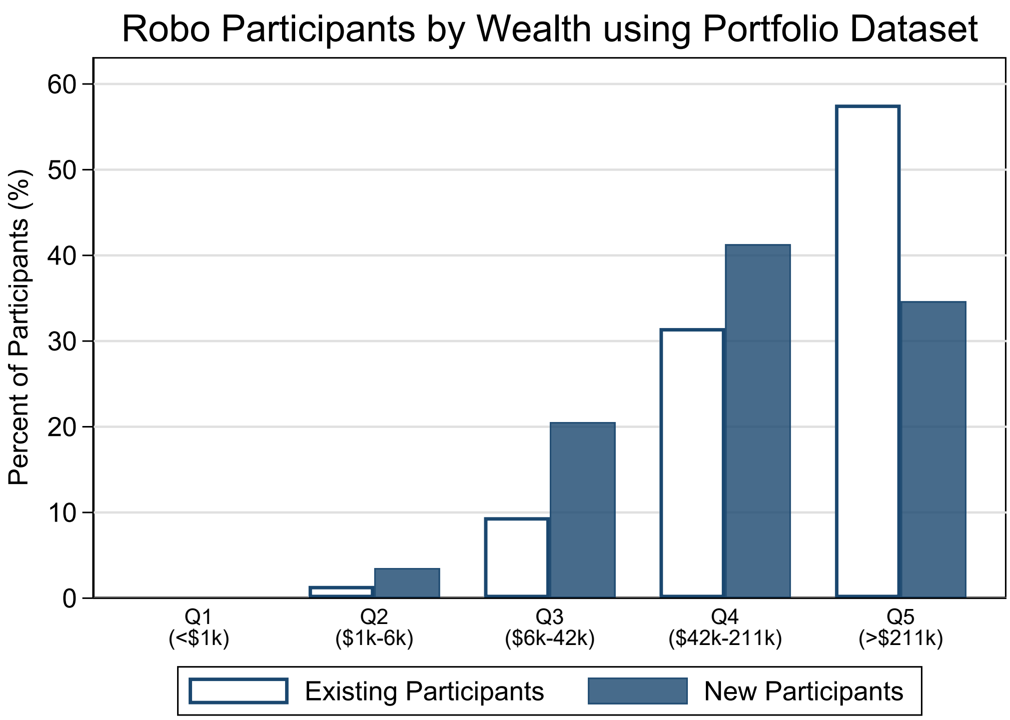

Note: This figure replicates Figure 3 using the subset of robo participants who are also in the portfolio dataset. The remaining notes are the same as in Figure 3.

Figure A4: Share of U.S. Households in the Middle Class by Age

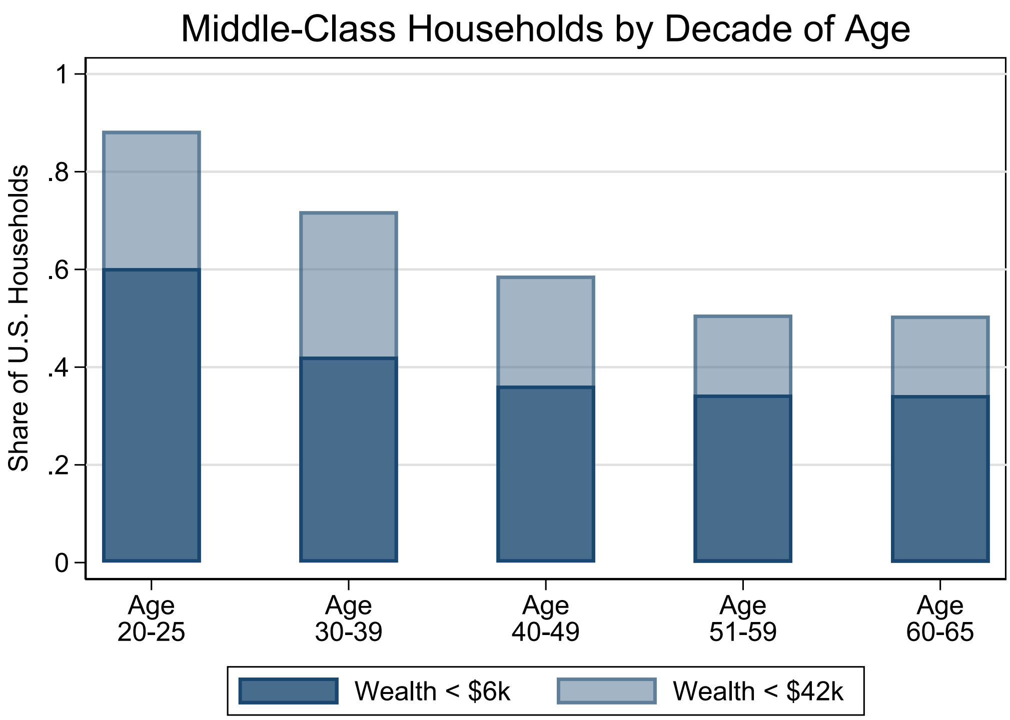

Note: This figure plots the share of U.S. households in the first or second quintile of the overall U.S. distribution of liquid assets (under $6k) and also the share in the first, second, or third quintiles (under $42k). The shares are calculated for each decade of investor age. Data are from the SCF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

Figure A5: Average Ad Spending Across Market Areas by Wealth

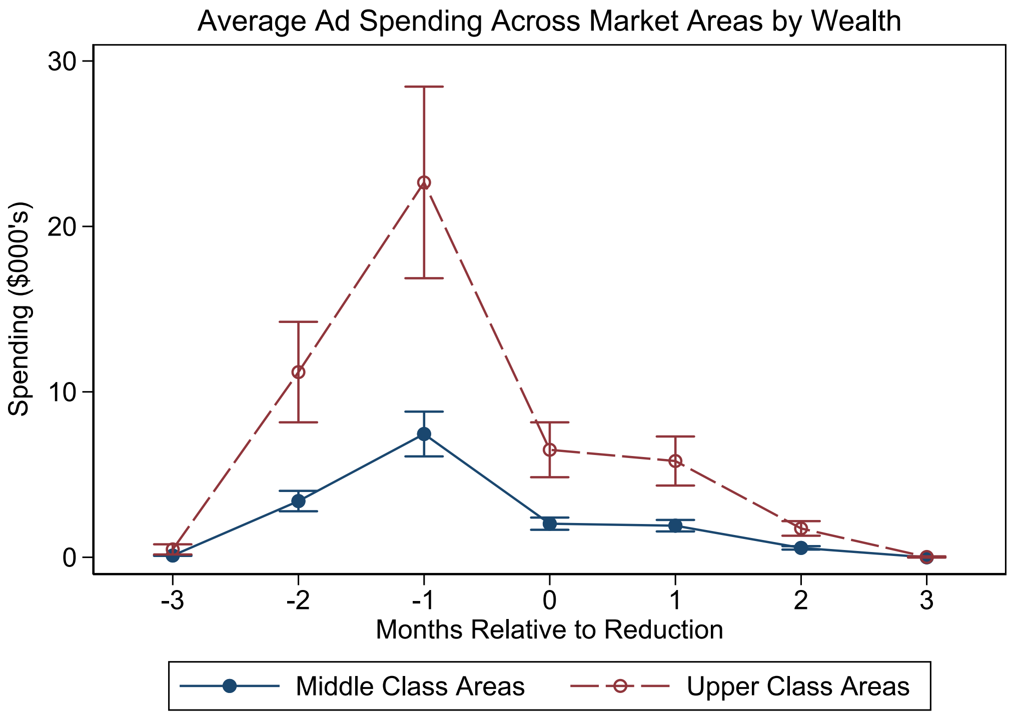

Note: This figure presents the times-series of average monthly ad spending by the advisor across middle-class and upper-class Designated Market Areas (DMAs). Time 0 is the month when the reduction in account minimum was implemented. A DMA is classified as middle-class if its average imputed stock market wealth falls within quintiles 2 and 3 of the wealth distribution across all DMAs. DMAs in the top two quintiles are classified as upper-class. The ad spending data are sourced from Vivvix by Kantar. Details on data construction are outlined in Appendix A. Brackets represent 95% confidence intervals.

Figure A6: Robo Value Exposure by Age and Wealth

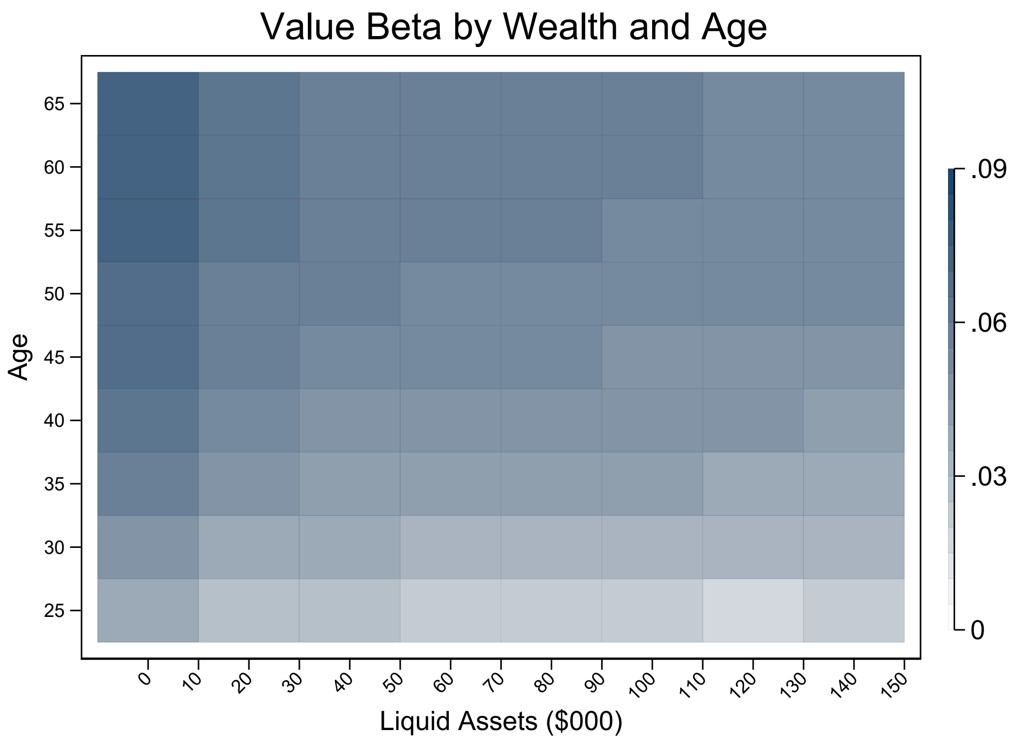

Note: This figure is analogous to Figure 1, and it shows the loading on the High-minus-Low factor (Value Beta). The remaining notes are the same as in Figure 1.

Figure A7: TDF Stock and Bond Exposure by Age and Wealth

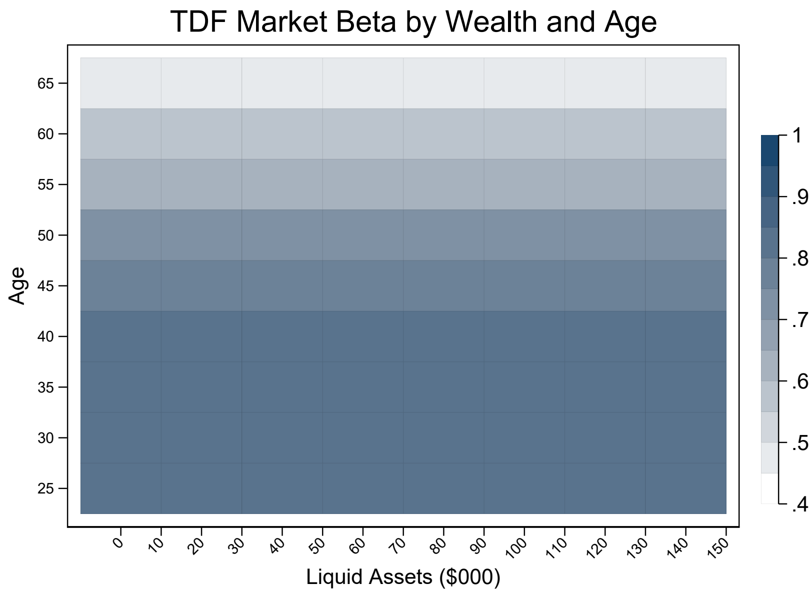

Panel (a)

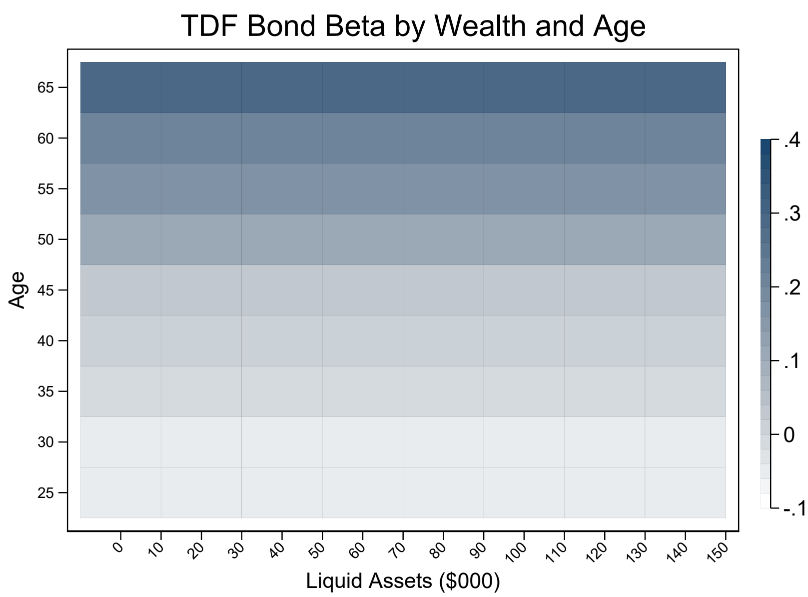

Panel (b)

Note: This figure is analogous to Figure 1, and it shows the loading on the market factor (*Market Beta*) and U.S. bond factor (*Bond Beta*) by investor age and wealth for the Fidelity Freedom TDF with target year that an investor of a given age would choose in 2015 to retire at age 65. The remaining notes are the same as in Figure 1.

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 Investors (%) (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 investors 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, VWO), 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: Primary and Secondary ETFs in Robo Portfolios

Asset Class	Primary ETF	Secondary ETF	Correlation Coefficient
U.S. Stocks	VTI	SCHB	0.999
High-Dividend Stocks	VIG	SCHD	0.948
Developed Market Stocks	VEA	SCHF	0.998
Emerging Market Stocks	VWO	IEMG	0.993
U.S. Corporate Bonds	LQD	.	.
Emerging Market Bonds	PCY	EMB	0.894
Municipal Bonds	MUB	TFI	0.929
Treasuries	SCHP	VTIP	0.843
Commodities	XLE	VDE	0.998
Real Estate	VNQ	SCHH	0.994

Note: This table reports the primary and secondary ETFs for each asset class defined by the robo advisor. The robo advisor will use the secondary ETF for accounts under $5,000 if the primary ETF is too expensive. Most of our results are based on the primary ETF and the approximation that the robo advisor can hold fractional shares. Table A11 assesses robustness of portfolio-level statistics to this approximation. The rightmost column reports the coefficient of correlation in total return between the primary and secondary ETFs. The correlation is calculated over the longest available time series for each ETF through December 2022. The robo advisor does not have a secondary ETF to LQD for U.S. Corporate Bonds.

Table A4: 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.

Table A5: Summary of Investors in Portfolio Dataset

	All ($N = 1,913$)			Robo Participants ($N = 860$)			Robo Non-Participants ($N = 1,053$)		
	Mean	Standard Deviation	Median	Mean	Standard Deviation	Median	Mean	Standard Deviation	Median
Liquid Assets$_i$ ('000)	484.24	770.1	200	447.29	739.97	184.5	514.4	792.94	200
Income$_i$ ('000)	166.91	131.07	130	170.44	131.67	140	164.02	130.57	125
Age$_i$	38.24	11.91	35	35.55	10.17	33	40.44	12.74	37
Self-Managed Sharpe$_i$	0.46	0.16	0.47	0.44	0.17	0.45	0.47	0.16	0.49
Middle$_i$	0.19	0.39		0.19	0.39		0.19	0.39	
Participant$_i$	0.45	0.5							

Note: This table summarizes investors in the portfolio dataset. The sample consists of investors who consult the robo advisor for a free portfolio review, of which 45% become participants and the remainder do not. The left, center, and right panels summarize all investors in the dataset, the subsample of robo participants, and the subsample of non-participants, respectively. The variable *Self-Managed Sharpe*$_i$ equals the Sharpe ratio of the investor's self-managed portfolio, as in Table 4. The remaining notes are the same as in Tables 1 and 4.

Table A6: Effect on Participation after Dropping Investors with Multiple Accounts.

	First Difference	Dynamic DiD
$Y_i =$	*New Participant$_i$*	*New Participant$_{i,t}$*
	(1)	(2)
$Middle_i$	12.320***	
	(1.413)	
$Middle_i \times Post_t$		0.646***
		(0.095)
Controls	Yes	
State FE	Yes	
Investor FE		Yes
Week FE		Yes
Controls $\times$ $Post_t$		Yes
State FE $\times$ $Post_t$		Yes
R-squared	0.069	0.007
Number of Observations	7,034	463,232

Note: Standard errors are in parentheses. Point estimates are multiplied by 100. This table re-estimates the main specifications from Table 2 after restricting the sample to account holders with only a single account with the robo advisor, which applies to 88% of account holders. This exercise assesses whether the estimated effects in Table 2 are driven by gambling motives induced by a lower account minimum. Account holders are identified by the user ID assigned by the robo advisor. The remaining notes are the same as in Table 2.

Table A7: Effects on Participation Within Middle Class

$Y_i =$	*New Participant$_i$*		
	(1)	(2)	(3)
Lower Middle$_i$	27.068***	16.313***	15.065***
	(3.058)	(3.046)	(3.004)
Upper Middle$_i$	21.328***	14.928***	13.588***
	(1.214)	(1.387)	(1.370)
Controls	No	Yes	Yes
State FE	No	No	Yes
R-squared	0.033	0.067	0.097
Number of Observations	9,349	9,349	9,349

Note: Standard errors are in parentheses. This table reestimates the effects from Table 2 separating between investors from the second U.S. wealth quintile (*Lower Middle$_i$*) and the third quintile (*Upper Middle$_i$*). We use a specification of the form:

$$New\ Participant_i = \mu_1 Lower\ Middle_i + \mu_2 Upper\ Middle_i + \psi X_i + \varrho + u_i.$$

The remaining notes are the same as in Table 2.

Table A8: Robustness of Democratization to State-Clustered Standard Errors

$Y_i =$	*New Participant$_i$*	
	(1)	(2)
Middle$_i$	15.051***	13.718***
	(0.901)	(1.070)
Controls	Yes	Yes
State FE	No	Yes
R-squared	0.067	0.097
Number of Observations	9,349	9,349

Note: Standard errors are in parentheses. This table reestimates the specifications in columns (1)-(2) of Table 2 after clustering standard errors by state of residence.

Table A9: 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 investors from the second or third U.S. wealth quintiles as constrained by the previous account minimum. Subscripts i and t index investor 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 investor 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 investor-weeks in which the investor makes its initial deposit, and, thus, we cannot include investor fixed effects. Columns (4)-(5) estimate separate effects for investors 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 A10: Interpreting the Magnitude of the 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.1. 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 A11: Summary of Self-Managed and Robo Portfolios without Fractional Shares

	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	0.748	0.296***	0.459	0.755	0.296***
Expected Return	0.080	0.101	0.022***	0.078	0.101	0.023***
Total Volatility	0.209	0.136	-0.072***	0.196	0.134	-0.062***
Idiosyncratic Volatility	0.146	0.0343	-0.113***	0.138	0.033	-0.105***
Robo-Sharpe Expected Return	0.154			0.147		
Return Loss	-0.074			-0.069		

Note: This table is analogous to Table 4 without approximating robo portfolio weights as though the advisor can hold fractional shares. The robo portfolio weights used in this table solve an optimization problem similar to that of the robo advisor (Wealthfront (2023)). Explicitly, we minimize the excess cash in the portfolio, subject to the constraints that: robo portfolios hold a discrete number of shares of any ETF; and the tracking error relative to the target weight for each asset class lies below some threshold. The solution requires selecting the ETF with the lowest price among the primary and secondary ETFs for each asset class shown in Table A3. We compute this ranking using each ETF's average price over the post-reduction period. We preserve computational feasibility by restricting the solution to equal either the integer floor or the integer ceiling of the fractional number of shares implied by the ETF's target weight. This restriction also implies a bound on the admissible tracking error. We solve this problem for all portfolios smaller than $5,000 because the fractional shares approximation is less accurate for small portfolios. The resulting average excess cash share equals 0.1% for the portfolios shown in the table, and it equals 0.2% for the subset of portfolios held by middle-class investors. Given the revised weights, we recalculate the portfolio statistics as in Table 4, assuming that the factor loadings for the primary and secondary ETFs in each asset class are the same, and similarly for the idiosyncratic volatilities. This assumption accords with the high correlation coefficients shown in Table A3. The remaining notes are the same as in Table 4.

Table A12: Summary of Self-Managed and Robo Portfolios excluding Possible Gambling Accounts

	Middle Class ($N = 351$)			Upper Class ($N = 1,497$)		
	Self-Managed	Matched Robo	Difference	Self-Managed	Matched Robo	Difference
	(1)	(2)	(3)	(4)	(5)	(6)
Sharpe Ratio	0.451	0.751	0.299***	0.458	0.754	0.296***
Expected Return	0.080	0.102	0.0230***	0.078	0.101	0.0230***
Total Volatility	0.209	0.137	-0.072***	0.196	0.134	-0.061***
Idiosyncratic Volatility	0.147	0.034	-0.112***	0.136	0.0330	-0.103***
Robo-Sharpe Expected Return	0.155			0.146		
Return Loss	-0.075			-0.068		

Note: This table is analogous to Table 4 after excluding self-managed accounts smaller than 10% of the investor's liquid wealth, which are possibly intended to serve as gambling accounts. The remaining notes are the same as in Table 4.

Table A13: Summary of Self-Managed and Robo Portfolio Loadings

	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)
Factor Loadings (β_i)						
Market	0.930	0.893	-0.036***	0.899	0.876	-0.023***
SMB	0.044	0.003	-0.040***	0.030	-0.001	-0.032***
HML	-0.086	0.061	0.147***	-0.074	0.061	0.135***
GLB	0.629	-0.020	-0.649***	0.574	-0.023	-0.597***
USB	-0.447	0.508	0.955***	-0.391	0.512	0.903***

Note: This table summarizes the factor loadings for self-managed and matched robo portfolios, based on the Fama-French Three Factor Model augmented with two bond factors. 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; USB is the return on the Bloomberg-Barclays Aggregate U.S. Bond Index Unhedged; and GLB is the return on the Bloomberg-Barclays Global Aggregate Bond Index Unhedged. The remaining notes are the same as in Table 4.

Table A14: Covariances and Risk Prices of 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): Risk Prices

	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 A13. The statistics are calculated as follows. The covariance matrix is that of: (i) the return on the CRSP Value-Weighted Index minus the one-month Treasury yield; (ii) the HML return for the U.S. market from Ken French's website; (iii) the SMB return for the U.S. market from Ken French's website; (iv) the return on the Bloomberg-Barclays Aggregate U.S. Bond Index Unhedged; and (v) the return on the Bloomberg-Barclays Global Aggregate Bond Index Unhedged. The covariance matrix is calculated using the subperiod from January 1990 through February 2016 because we only observe the bond factors over that subperiod. We calculate risk prices using the sample means of the series (i)-(v), taken over January 1960 through December 2017 for (i)-(iii) and over January 1990 through February 2016 for (iv)-(v). 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 A13.

Table A15: 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***
CAPM	0.334	0.442	0.108***	0.339	0.439	0.100***
(b) Expected Return						
Fama-French	0.064	0.071	0.007***	0.062	0.069	0.007***
CAPM	0.058	0.063	0.006***	0.056	0.062	0.006***
(c) Total Volatility						
Fama-French	0.213	0.137	-0.077***	0.197	0.134	-0.063***
CAPM	0.212	0.143	-0.068***	0.198	0.142	-0.057***
(d) Idiosyncratic Volatility						
Fama-French	0.151	0.036	-0.115***	0.142	0.035	-0.106***
CAPM	0.173	0.079	-0.094***	0.160	0.079	-0.081***
(e) Return Loss from Difference in Sharpe Ratio						
Fama-French	-0.046			-0.039		
CAPM	-0.035			-0.030		

Note: This table assesses the robustness of Table 4 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 4.

Table A16: Additional Heterogeneity in Welfare Gain

	Increase in Lifetime Consumption
(a) Extensive vs. Intensive Margin	
New Participants	0.77%
Existing Participants	0.05%
Weighted Average	0.62%
(b) New Participants by Initial Deposit	
Initial Deposit: $500 to $2,000	0.86%
Initial Deposit: $2,000 to $3,500	0.70%
Initial Deposit: $3,500 to $4,500	0.55%
(c) New Participants by Wealth Quintile	
Quintile 2	1.21%
Quintile 3	0.52%

Note: This table calculates the welfare gains from the reduction for various subpopulations of robo participants. Panel (a) reports the gain separately for investors who become participants after the reduction (New) and who participated under the previous minimum (Existing). Note that Table 7 only reports the welfare gain for new participants. The bottom row of panel (a) shows the weighted average welfare gain across new and existing participants. Panel (b) reports the gain for new participants according to whether their initial deposit lies in one of the indicated ranges. Panel (c) reports the gain for new participants according to whether they are in the second or third U.S. wealth quintile. All panels restrict to middle-class participants. The remaining notes are the same as in Table 7.

Table A17: Summary of Robo Participants by Wealth Quintile.

	Existing Participants			New Participants			
	Mean	Standard Deviation	Median	Mean	Standard Deviation	Median	Difference in Mean
(a) Third Wealth Quintile:							
Initial Deposit$_i$ ('000)	7.74	5.51	5	5.35	13.33	2	-2.387***
Liquid Assets$_i$ ('000)	25.43	10.41	25	21.83	10.5	20	-3.6***
Income$_i$ ('000)	91.23	60.69	80	69.21	43.26	61.27	-22.015***
Age$_i$	30.27	6.39	29	30.06	6.96	29	-0.215
			$N = 568$			$N = 1,415$	
(b) Second Wealth Quintile:							
Initial Deposit$_i$ ('000)	6.05	2.31	5.02	2.1	3.33	1	-3.945***
Liquid Assets$_i$ ('000)	5.15	.37	5	4.53	.95	5	-0.623***
Income$_i$ ('000)	106.31	72.65	90	52.34	33.29	50	-53.967***
Age$_i$	30.82	5.77	30	29.91	7.83	28	-0.910
			$N = 69$			$N = 198$	

Note: This table summarizes a subset of variables from Table 1 for various subsamples defined by the investor's wealth. The remaining notes are the same as in Table 1.

Table A18: 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%
(b) Calibration of Portfolio Parameters				
Calvet, Campbell and Sodini (2007)	0.88%	0.64%	0.66%	1.97%
(c) Calibration of Other Parameters				
Larger Reduction (\$5,000 to \$200)				
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%
Procyclical Labor Income ($\rho^Y = 80\%$)	0.72%	0.51%	0.58%	1.61%
(d) Other Counterfactuals				
Comparison portfolio is Fidelity TDF	0.16%	0.05%	0.10%	0.57%

Note: This table recalculates the welfare gains from the reduction under various extensions and reparameterizations of the baseline model in Table 7. 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); and 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). The participation cost κ equals either \$200, based on the GMM estimates in Appendix Table A19, or a higher value of \$700. Panel (b) assesses robustness to the calibration of the self-managed portfolio parameters β_i^S and $\sigma_{\epsilon,i}$, where we modify the projection of portfolio characteristics onto the model's state variables described in Section 6.2.1. 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 correlation coefficient between log labor income and financial returns of $\rho^Y = 80\%$ corresponds to a loading of β^Y of 3.3. Panel (d) computes the gain under a counterfactual where the self-managed match is replaced by the Fidelity Freedom Fund TDF in which an investor of a given age t in 2015 and retirement age 65 would invest, not the Vanguard TDF as used in Table 7. The remaining notes are the same as in Table 7.

Table A19: GMM Estimates of Preference Parameters

Parameter	Estimate
Coefficient of Relative Risk Aversion (γ)	5.1
	[4.420, 5.830]
Discount Factor (δ)	0.92
	[0.880, 0.962]
Stock Market Participation Cost (κ)	168.75
	[31.338, 306.162]

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. Each empirical moment is weighted by its inverse standard error (Newey (2007)). 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. Standard errors are bootstrapped. Brackets correspond to 95% confidence intervals under an asymptotically normal distribution.