Automation and Inequality in Wealth Management*

Michael Reher† and Stanislav Sokolinski‡

March 2022

Abstract

We examine how access to automated wealth managers affects households' investment in financial markets and welfare across the wealth distribution. Our setting features novel microdata from a major U.S. robo advisor and a quasi-experiment in which the advisor reduces its account minimum by 90%. Based on a difference-in-difference estimator, the reduction increases middle-class households' participation by 110% but does not affect wealthier or poorer households. We rationalize this behavior with a life cycle model calibrated using portfolio-level data. Our calibration suggests that the reduction significantly raises middle-class households' welfare, and 65% of this gain reflects improved diversification.

Keywords: FinTech, Financial Advice, Portfolio Delegation, Inequality
JEL Classification: G11, G24, D3, O3

*We thank Christoph Aymanns, John Campbell, Francesco D'Acunto, Ben Friedman, Pedro Gete, Steffen Hitzemann, Jakub Jurek, Iva Kalcheva, Yaron Levi, Jun Liu, Matteo Maggiori, Apurva Mehta, Will Mullins, Alessandro Previtero, Franklin Qian, Enrichetta Ravina, Jonathan Reuter, Alberto Rossi, Andrei Shleifer, Celine Sun, Ross Valkanov, Boris Vallee, anonymous referees, and seminar participants at the University of Wisconsin, Jackson Hole Winter Finance Conference, FINRA NORC Diversity Conference, FMA Annual Conference, Paris FinTech Conference, Toronto FinTech Conference, the New York Fed's FinTech Conference, Georgetown's Finance and Policy Center, UC San Diego, Rutgers Business School, California Corporate Finance Conference, CAFR FinTech Workshop, and Harvard's finance PhD lunch for comments. Ruikai Ji provided excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the position of Wealthfront Inc.

†University of California San Diego, Rady School of Management. Email: mreher@ucsd.edu
‡Rutgers Business School, Newark and New Brunswick. Email: ssokolinski@business.rutgers.edu

1 Introduction

Wealth managers oversee $37 trillion in household assets, but most of these managers cater only to the affluent.[1] This observation undergirds a well-known argument that unequal access to investment opportunities perpetuates wealth inequality (e.g., Piketty 2014). Implicit in this argument is the untested assumption that modestly wealthy households would benefit from professional wealth management if only they, too, could access it. This question has great practical relevance given the tenfold growth of automated wealth managers (i.e., robo advisors) over the past decade, many of which promote the idea of democratizing wealth management.[2] A budding literature studies how robo advisors affect the relatively wealthy, who already have access to financial professionals (e.g., D'Acunto, Prabhala and Rossi 2019; Loos et al. 2020; Rossi and Utkus 2020). Building on this literature, we use robo advisors as an example from which to learn how access to wealth management affects a much larger, less-wealthy population.

We show that the modestly wealthy can significantly benefit from access to automated wealth management. First, using novel microdata and a quasi-experiment, we find that households from the lower-middle quintiles of the U.S. wealth distribution become significantly more likely to participate in automated wealth management when it becomes more accessible, whereas wealthier households show no response. Then, we quantitatively explain this finding with a life cycle model in which households can invest independently or through a professional. Calibrating the model to portfolio-level data, we show that middle-class households optimally participate in wealth management when it becomes more accessible because they struggle to diversify on their own. In terms of magnitude, the reduction improves middle-class households' welfare by the same amount as a 4 percentage points (pps) increase in the equity premium. Surprisingly, those over age 55 gain the most. Our results suggest that automation, in the form of robo advice, leads to Pareto improvements by reducing inequality in wealth management.

Identifying the effect of access to wealth management on the modestly wealthy is challenging for two reasons. The first challenge is acquiring data. Regulatory filings, industry reports, and other public datasets do not contain information about the composition of households who participate with specific wealth managers, which is central to our paper. Accordingly, we obtain

[1] Wealth managers' retail assets equaled $37 trillion in 2020 (Heredia et al. 2020). Opening an account with a private wealth manager typically requires a minimum investment of at least $100,000 (Pilon 2011).

[2] Quoting the financial press: "The wealth-management industry stratifies customers in a manner rather similar to airlines. High-net-worth clients fly business class, picking stocks and chatting in person with named advisors. Cattle class gets no service at all. Technology is conspiring to change that" (The Economist 2019). The top five robo advisors managed $283 billion in 2020 versus $30.4 billion in 2015 (Appendix Table A1).

two novel datasets directly from a major U.S. robo advisor. In one dataset, we observe the demographic background, investment activity, and liquid assets (including retirement accounts), of households who participate with the advisor. This information enables us to empirically analyze the distributional effects of automated wealth management. In the second dataset, we observe pairs of portfolios for households interested in becoming robo participants. Each pair contains the non-robo portfolio they manage themselves and the robo portfolio they would receive. We use this information to structurally decompose the channels through which automated wealth management improves welfare.

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

We find that the reduction democratizes the market for automated wealth management. The wealth distribution of participants shifts sharply leftward after the reduction, while showing no pre-trend in the months leading up to it. In particular, the share of participants from the second and third U.S. wealth quintiles, whom we call the "middle class", increases by 107% (16 pps).[3] This increase reflects a sharp break from trend that is not present among participants from the upper two quintiles, whom we call the "upper class". However, the democratization is asymmetric, in that there is no change in participation among the poorest quintile.

We formalize this graphical intuition through a difference-in-difference analysis. Our regression model compares the probability of participating with the robo advisor after versus before the reduction between middle versus upper-class households. Intuitively, the middle class represents the "treated" group in that it experiences a relaxation of minimum-account constraints due to the reduction. Accordingly, we find that middle-class households are 14 pps more likely to participate

[3]Most Americans would use the term "lower-middle class" to describe what we call the "middle class" and the term "upper-middle class" to describe what we call the "upper class" (Reeves (2015)).

with the robo advisor after the reduction, relative to the upper class. We show that this estimate implies a 110% increase in the total number of middle-class robo participants. This finding is robust to various measures of wealth and, thus, is not biased by measurement error.

To check that we identify the desired effect, we examine whether a relaxation of minimum-account constraints drives the results. Otherwise, the middle class could have accessed wealth management without the reduction. Consistent with the constraints channel, the majority of middle-class households who became robo participants prior to the reduction bunched their investment at the previous minimum of $5,000, a hallmark of binding constraints (e.g., DeFusco, Johnson and Mondragon (2020)). After the reduction, however, such bunching immediately disappears, and most new middle-class participants make a previously infeasible investment of under $5,000. Indeed, based on a wide variety of tests, we find no evidence that the results are driven by channels distinct from a relaxation of constraints, such as: heightened visibility from targeted advertising or media attention (e.g., Kaniel and Parham (2017)); gambling motives (e.g., Bombardini and Trebbi (2012)); business stealing from competitors; or heterogeneous trends by demographics or risk attitude.

Our empirical findings raise both positive and normative questions, neither of which can be answered through a purely reduced-form analysis. In positive terms, it is not obvious how to rationalize the reduction's large effect in light of existing theories. On the one hand, workhorse quantitative models assume households can invest efficiently on their own and, thus, do not need professional management (e.g., Cocco, Gomes and Maenhout (2005)). On the other hand, recent analytic models show how households may optimally demand professional management if they cannot invest well on their own because of, say, underlying anxiety or limited information (e.g., Gennaioli, Shleifer and Vishny (2015); Gârleanu and Pedersen (2018)). Quantitatively, however, analytic models do not inform whether the drawbacks of self-management could reasonably generate such large demand as our estimates imply. In normative terms, it is not obvious whether access to wealth management meaningfully improves welfare and, if so, through what channel. For example, robo portfolios may feature a higher Sharpe ratio than investing on one's own (e.g., D'Acunto and Rossi (2020)), but they may yield lower utility if, say, the higher Sharpe ratio reflects more exposure to priced risk than households want to take.

We address these questions by adding two novel ingredients to a workhorse life cycle model. First, households choose between managing their own portfolio or delegating it to a professional. We do not make any assumptions about the differences between self-managed and delegated (i.e.,

robo) portfolios. Instead, we use our portfolio-level dataset to calibrate each portfolio's risk profile across the joint distribution of household age and wealth. Accordingly, we find that the two portfolios differ in their quantity of priced versus idiosyncratic risk. In particular, robo portfolios have a 2 pps higher expected return than self-managed ones, which comes from exposure to priced stock and bond risk factors. However, the total risk in robo portfolios is much lower, reflecting an 11 pps lower idiosyncratic volatility. The model's second distinguishing ingredient is an account minimum required by professional managers, which makes professional management less accessible. To reiterate, we add these two realistic ingredients to a workhorse model and ask whether it can qualitatively and quantitatively explain our empirical results.

The model formalizes the following intuition. Suppose households have limited ability to independently diversify idiosyncratic risk (e.g., Calvet, Campbell and Sodini (2007); Von Gaudecker (2015)) or take priced risk (e.g., Gennaioli, Shleifer and Vishny (2015); Hitzemann, Sokolinski and Tai (2021)). Given this limited ability, they optimally allocate a share of their wealth to professional managers. However, middle-class households cannot achieve this optimal share because it requires an investment below the required minimum. Therefore, they can either invest more than their optimal share or simply not participate in wealth management. In the latter case, reducing the minimum relaxes minimum-account constraints and so prompts them to participate. The poorest households, however, may still find that participating in wealth management requires an excessive risky share. Consequently, the reduction has an asymmetric effect on participation across the wealth distribution, as we find empirically.

We find strong support for this intuition, in that the model quantitatively replicates our key empirical findings on both the extensive (e.g., growth in participation) and intensive margins (e.g., portfolio share). Thus, from a positive perspective, we interpret the large estimated effect of the reduction as an optimal response to differences in portfolio characteristics. Importantly, our framework does not require us to explain why, for example, households do not diversify on their own. Rather, in the spirit of a sufficient statistics approach, we take these differences as given. In particular, we do not need to specify the magnitude of difficult-to-quantify behavioral parameters, such as peace-of-mind (e.g., Gennaioli, Shleifer and Vishny (2015)) or costs of stock market participation (Vissing-Jørgensen (2003)).

Turning to welfare, the reduction raises middle-class households' welfare by 2%, based on the standard lifetime consumption metric, and it has almost no effect on the upper class. For reference, a permanent 4 pps increase in the equity premium (from 7.6% to 11.6%) with no reduction would

also raise welfare by roughly 2%. Thus, households value access to robo portfolios under the reduced minimum as much as they value a large (4 pps) increase in the equity premium with no such access. We structurally decompose the channels through which the reduction improves welfare, finding that 65% reflects reduced idiosyncratic volatility, 15% reflects greater priced risk exposure, and 20% reflects a higher risky share in response to the previous two gains. Thus, automated wealth management adds value principally through diversification. We find similar welfare results under various parameterizations and under extensions in which households can borrow or invest in a defined contribution plan.

Lastly, middle-class households over age 55 gain 0.4 pps (22%) more in welfare than those under age 36. Intuitively, the reduction relaxes a permanent constraint on older middle-class households but only a temporary constraint on younger ones. In particular, our model implies that 83% of new participants over age 55 would never have participated in wealth management without the reduction, being permanently constrained. By contrast, 77% of new participants under age 36 would have eventually overcome the previous minimum, being temporarily constrained. Therefore, robo advisors add significant value to households at a later stage in life, despite their popular image as the investment of choice for millennials.

From a policy perspective, a number of government programs have attempted to expand the set of investment opportunities available to modestly wealthy households, with mixed rates of success (e.g., myRA, OregonSaves, NEST). Our results exemplify how private, automated wealth management can itself improve the financial condition of the modestly wealthy. From the perspective of economic theory, our results support models of bounded rationality in which households act optimally given limits on their ability to invest efficiently on their own.

We conclude this section by situating our contribution within the literature. Section 2 provides institutional background and describes our quasi-experiment. Section 3 describes our data. Section 4 estimates the effect of the reduction on the democratization of the robo market, and Section 5 assesses its robustness. Section 6 describes the life cycle model. Section 7 studies positive implications. Section 8 studies welfare implications. Section 9 concludes. The online appendix contains additional material.

Related Literature

We principally contribute to two strands of literature. First, we contribute to a literature on new financial technologies (FinTech) by showing how they democratize wealth management. Within

this literature, we most directly contribute to an emerging agenda on robo advisors that, to date, has examined how they affect wealthier households. We differ not only in our focus on the modestly wealthy, but also in that our setting features: full portfolio delegation, as opposed to non-binding suggestions (D'Acunto, Prabhala and Rossi (2019); Bianchi and Briére (2020); D'Hondt et al. (2020)); no option for human advice (Rossi and Utkus (2020)); robo advice unaffiliated with the banking system (Loos et al. (2020)); and quasi-experimental evidence (Reher and Sun (2019)). Methodologically, to the best of our knowledge we are the first to structurally analyze the effects of robo advice. This approach enables us to evaluate long-term welfare gains from automated wealth management, complementing evidence on short-term changes in asset allocation.

More broadly within the FinTech literature, we show how FinTech affects financial inclusion and wealth inequality in a new setting. This finding complements analogous results in the contexts of app-based payments (e.g., Hong, Lu and Pan (2020)), bank deposits (e.g., Bachas et al. (2018); Bachas et al. (2020); Higgins (2020)), and mortgage markets (e.g., Fuster et al. (2019); Bartlett et al. (2021); Fuster et al. (2021)). In terms of inequality, our empirical results confirm the theoretical prediction of Philippon (2019) that robo advising favors the middle class over both the upper and lower classes.[4]

Second, we contribute to the household finance literature theoretically and empirically. Theoretically, we show how households optimally seek professional portfolio management when they cannot diversify on their own. In so doing, we follow in a long tradition of quantitative life cycle models summarized by Gomes (2020). Our model stands out in that we match quasi-experimental evidence, incorporate both self-managed and professionally managed portfolios, and, like Fagereng, Gottlieb and Guiso (2017), calibrate it using microdata. These features allow us to parsimoniously match the data, without, for example, the reduced-form costs of stock market participation (e.g., Vissing-Jørgensen (2003)) that comparison models often require.

Empirically, we contribute to the household finance literature by characterizing account minimums as a novel friction that constrains investment in risky asset markets. This friction arises from the supply side and does not directly depend on household characteristics such as preferences (e.g., Barberis, Huang and Thaler (2006)), sophistication (e.g., Grinblatt, Keloharju and Linnainmaa (2011); Christelis, Jappelli and Padula (2010)), socialization (e.g., Hong, Kubik and

[4]It is well-known that financial returns increase in wealth and education, and our findings suggest that FinTech can reduce this inequality (e.g., Lusardi, Michaud and Mitchell (2017); Campbell, Ramadorai and Ranish (2019); Bach, Calvet and Sodini (2020); Fagereng et al. (2020)). This conclusion differs from that of models predicting a positive relationship between FinTech and inequality (e.g., Kacperczyk, Nosal and Stevens 2019; Mihet 2020).

Stein (2004)), or education (e.g., Cole, Paulson and Shastry (2014); Van Rooij, Lusardi and Alessie (2011)). Our findings caveat Haliassos and Bertaut (1995) who, based on a two-period model, conclude that account minimums at discount brokers do not affect household investment.

2 Institutional Background

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

2.1 The U.S. Robo Advising Market

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

At the time of our analysis, Wealthfront managed roughly $3 billion and was the largest standalone robo advisor in the U.S. market, with Betterment and Personal Capital as its nearest competitors. Two traditional managers, Vanguard and Charles Schwab, launched robo advising services early in 2015. Both of these services managed more than Wealthfront because they transferred assets from existing, non-robo services. 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. Wealthfront is the only robo advisor that relies purely on automation, with no option for a human advisor.

2.2 The Robo Advisor

Wealthfront, henceforth "the robo advisor", has offered many services throughout its history, including tax loss harvesting, long term financial planning, portfolio lines of credit, and a risk parity fund. Most relevant for this paper is its baseline product: an automatically rebalanced portfolio of 10 ETFs corresponding to 10 asset classes.[5] The portfolio weights are determined by a questionnaire that asks the client several questions about age, liquid assets, income, demographic background, and response to hypothetical investment decisions. The client is then assigned to one of 20 possible risk tolerance scores, which range from 0.5 to 10 in increments of 0.5. Each risk tolerance score uniquely determines a robo portfolio. The portfolio weights solve a problem of optimal asset allocation across the 10 ETFs, taking this score as a parameter. As summarized in Appendix Table A2, portfolios associated with higher risk tolerance scores exhibit higher betas, higher expected returns, and higher proportions of wealth invested in stocks.

The robo portfolios that we study conform to most "textbook" recommendations for retail investors (e.g., Malkiel 2015). They provide well-diversified risk exposure with more personalization than a generic "60/40" portfolio, but without the complexity often associated with active management. Importantly, robo portfolios are not recommendations, but, rather, they are directly managed by the robo advisor. Consequently, households have little discretion over their portfolio allocations, and so their robo performance will not depend on sophistication (e.g., Grinblatt, Keloharju and Linnainmaa (2011); Christelis, Jappelli and Padula (2010)), ability to diversify (e.g., Calvet, 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)).

2.3 The 2015 Reduction in Account Minimum

On July 7, 2015, the robo advisor unexpectedly reduced its account minimum from $5,000 to $500, meaning that a household would need to invest only $500 to participate with the advisor as opposed to $5,000 beforehand. This reduction is quite sizeable from the standpoint of most U.S. households. For reference, $5,000 equals 30% of the median household's liquid assets ($17,000), and it defines the 37th percentile of the U.S. wealth distribution, according to the 2016 Survey of Consumer Finances. Prior to the reduction, therefore, half of U.S. households could not

[5] Strictly speaking, each asset class has a primary ETF and multiple secondary ETFs. The robo advisor will rebalance toward the secondary ETF if doing so yields a capital loss and, thus, reduces the client's tax liability. The 10 primary ETFs are chosen to track stock market indices (VIG, VTI, VEA, VW), bond market indices (LQD, EMB, MUB, TIPS), and other asset classes, namely real estate (VNQ) and commodities (XLE).

participate without investing at least 30% of their wealth, while 37% could not participate at all without borrowing. The reduction was motivated by the advisor's philosophy of inclusive investment and belief that non-wealthy households will eventually accumulate enough assets to become high-revenue customers.[6] Indeed, given the advisor's management fee of 0% for accounts under $10,000 (0.25% for larger accounts), the reduction was not intended to raise short-term revenue.

At the time of the reduction, all of the largest five U.S. robo advisors required an account minimum of at least $5,000 except for one, which had no account minimum but maintained a fee structure that discouraged setting up small accounts.[7] Importantly, the month of the reduction does not coincide with any other product launches by the robo advisor, any changes in its fee, or any significant developments in the overall robo advising market. This effectively idiosyncratic timing allows us to identify the reduction's effect on household participation in automated wealth management in Section 4.

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

3 Data

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

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

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

3.1 Deposits Dataset

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

Studying a company-specific dataset has two advantages over publicly available datasets such as the SEC's Form ADV filings, which serve as the basis for many industry reports about the robo market. First, estimates of robo participation growth derived from Form ADV data would be highly imprecise because they can include both inactive clients and "clients" who create a username but never provide the robo advisor with any money.[8] Second, unlike public data, our dataset includes information about a robo participant's wealth, allowing us to study investment activity across the wealth distribution.

3.2 Portfolio Dataset

The second core dataset contains snapshots of self-managed, non-robo portfolio holdings for both robo participants and robo non-participants. The advisor obtains these snapshots from an online tool through which it provides free financial advice to candidate clients about their outside portfolio holdings. As with the deposits dataset, we obtained this portfolio dataset through a direct query of the robo advisor's internal server. We merge this dataset with security-level information from standard sources (e.g., CRSP) to produce a cross-section of 1,913 portfolio pairs. Each pair consists of a self-managed portfolio and a counterfactual robo portfolio that the candidate

[8]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).

client would receive by becoming a robo participant.

We use the portfolio dataset to calibrate the life cycle model in Section 6, and two features make the dataset ideal for this purpose. First, we observe overall advisory and management fees. By restricting the set of non-robo portfolios to those without such fees, we can ensure with a high degree of confidence that the non-robo portfolios we study are managed by their owner (i.e., self-managed). Therefore, we can realistically parameterize the choice between managing one's own portfolio versus delegating it to a professional. Second, we observe the portfolios not only of robo participants (45% of sample), but also of robo non-participants (55% of sample). Observing non-participants helps our calibration avoid selection bias from, say, the possibility that only households who cannot invest efficiently on their own delegate to a professional. Importantly, selection from the fact that the dataset only includes households who consult the online tool does not pose a concern. These households are on the margin of robo participation, and, thus, they are exactly those whose behavior we seek to model.

In sum, each dataset accomplishes a separate goal. The deposits dataset enables our main empirical analysis: to examine how expanded access to automated wealth management affects the wealth distribution of robo participants. The portfolio dataset enables us to calibrate the model that explains our empirical results, since it includes information on robo and non-robo portfolios for both robo participants and non-participants.

3.3 Auxiliary Datasets

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

3.4 Summary Statistics

Table 1 summarizes the deposits dataset, the key dataset for our empirical analysis.[9] Panel (a) compares households who become robo participants after the reduction in account minimum (i.e., new participants) with existing participants. Reading from left to right, new participants are significantly less wealthy, earn lower incomes, make smaller initial deposits, and are 16 pps more likely to belong to the middle class, relative to existing participants. They are also more likely to come from less financially developed U.S. states, as shown in Appendix Figure A1.

Panel (b) conveys a similar pattern when restricting the comparison to middle-class households. In particular, the median new middle-class participant's initial deposit of $2,000 would have been infeasible under the previous account minimum of $5,000. Indeed, over half of existing middle-class participants invested exactly $5,000 for their initial deposit, suggesting that they were constrained by the previous account minimum. Interestingly, new middle-class participants are not significantly younger than existing ones, suggesting that the reduction does not work through generation-specific effects (e.g., technological savviness). We revisit this observation in Sections 4.3 and 8.2.

Lastly, new middle-class robo participants exhibit persistent investment behavior. For example, 97% do not close their account over our sample period, mirroring the 98% non-closure rate among all new participants. Additionally, 72% of new middle-class robo participants make a subsequent deposit, comparable to 71% among all new participants. Such subsequent deposit-making resembles the "dollar cost averaging" strategy commonly advocated by practitioners, which Brennan, Li and Torous (2005) show is optimal for risk averse investors.

4 Democratization of the Robo Market

We examine how the reduction in minimum increases robo participation by constrained, middle-class households, thereby democratizing the market for automated wealth management. First, we provide graphical evidence (4.1). Then, we formalize our identification strategy (4.2), report our main results (4.3), and assess the magnitude of the effect (4.4).

4.1 Graphical Evidence

Four pieces of graphical evidence suggest that reduction democratizes the robo market by

[9] We summarize the portfolio dataset later in Section 6 when we describe the model.

relaxing minimum-account constraints on the middle class. First, Figure 1 shows how the wealth distribution of robo participants shifts left after the reduction. This shift reflects how new robo participants are significantly less wealthy than existing ones, as already documented in Table 1.

Second, Figure 2 shows how the leftward shift documented in Figure 1 makes the robo wealth distribution more representative of the overall U.S. wealth distribution (i.e., more "democratic"). Notably, the share of robo participants from the second and third quintiles of the U.S. wealth distribution grows by 107% (16 pps), while the share from the upper two quintiles falls by 18% (16 pps). However, there is a non-monotonic relationship between robo participation growth and wealth, since lower-class households remain non-participants.

Third, Figure 3 shows how the increase in middle-class households' robo participation occurs strikingly and immediately after the reduction. In particular, the sharp jump and absence of a pre-trend in middle-class participation strongly suggests that this increase does not reflect reverse causality. Otherwise, an exogenous shock to middle-class robo participation coinciding exactly with the month of the reduction would have prompted the advisor to reduce its minimum at exactly that time, which seems implausible. More likely, the advisor accurately judged that reducing its minimum would induce such an increase in middle-class participation.

Fourth, Figure 4 shows how middle-class robo participants invest in a way consistent with binding constraints imposed by the previous minimum. Panel (a) shows how 65% of new middle-class robo participants invest under $5,000 after the reduction. The previous minimum would have precluded such a small investment. This behavior suggests that many middle-class households would have preferred to invest under $5,000 before the reduction, but they were constrained. Indeed, panel (b) shows that 52% of middle-class households who became participants before the reduction invest right at the minimum, a hallmark of constrained behavior. However, this bunching behavior dissipates after the reduction, consistent with a relaxation of constraints. Notably, these patterns are much less pronounced among upper-class households. This supports the idea that the change in participation between the middle versus upper classes represents the effect of minimum-account constraints, rather than other time-varying factors.

Collectively, the graphical evidence shows that a leftward shift in the robo wealth distribution occurs immediately after the reduction, making the distribution more representative of the U.S. population. In the remainder of this section, we formally test whether the reduction causally induces this shift by relaxing constraints on middle-class households.

4.2 Identification

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

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

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

We propose that the reduction affects robo participation among households with moderate levels of wealth because it relaxes constraints on their ability to invest. Equation (1) measures "moderate wealth" using the indicator $Middle_i$. Therefore, under an identification assumption described shortly, the parameter μ equals the effect of the reduction on middle-class households' probability of robo participation.[10]

The additional terms in equation (1) capture channels distinct from the minimum-account constraints channel. The fixed effect ζ_i captures slow-moving characteristics that predispose households to participating with the advisor, such as sophistication or trust (e.g., Guiso, Sapienza and Zingales 2008). Since such "affinity" to the advisor increases the probability of participation in any period, we can separately identify the effect of minimum-account constraints because the minimum changes over time. The interaction between X_i and $Post_{\mathcal{T}}$ captures heterogeneous trends by observed household characteristics. If, for example, younger households are more likely to become robo participants after the reduction for reasons apart from a relaxation of minimum-account constraints, then the coefficient ψ would separately capture this effect.

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

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

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

[10] Explicitly, μ equals the double difference in the probability of becoming a robo participant after versus before the reduction between middle-class versus upper-class households.

and $u_i \equiv \Delta v_i$. We estimate equation (2) on the set of eventual robo participants, and, therefore, μ equals the reduction's effect on the probability of robo participation, conditional on eventually participating.

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

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

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

Apart from measurement error in self-reported liquid assets, which we discuss at length below, there are two other ways in which equation (3) could be violated. First, u_i may capture changes in middle-class households' robo participation that coincide with the reduction, but which are not caused by it. One such confounding change could be trend growth in middle-class households' robo participation. However, the strong parallel trends shown in Figure 3 make this an unlikely source of bias. Another potentially confounding factor could be contemporaneous developments in the robo industry, such as the launching of new robo products by the two traditional managers named in Section 2.2. However, these new products were not targeted toward the middle class, and they were launched at least two months prior to the reduction, comfortably before the strong divergence in middle-class households' behavior in Figure 3.

Second, equation (3) could be violated if the reduction actually causes other shocks that affect middle-class households' robo participation. The leading examples are media attention and advertising. If middle-class households are more exposed to such media and advertising, then μ confounds the effect of heightened visibility with the effect of minimum-account constraints (e.g., Kaniel and Parham 2017). In Section 5.2, we test for bias from heightened visibility and find no evidence of it.

4.3 Baseline Results

Table 2 reports the results. The estimate in column (1) implies that middle-class households are 22 pps more likely to become robo participants after the reduction in account minimum, relative

to upper-class households. After we add household-level control variables (column (2)) and state fixed effects (column (3)), the estimate equals 14 pps, which we take as our baseline. Column (4) shows that the effect of the reduction does not vary with age, in line with the descriptive evidence in Table 1. This finding does not contradict the standard recommendation that younger households invest more in risky assets (e.g., Campbell and Viceira (2002)), since robo portfolios automatically become less risky as households age (e.g., Appendix Table A2). We also find that the effects of the reduction are stronger for risk-averse households, defined as those who request a less risky portfolio than that initially recommended by the advisor (column (5)). Intuitively, risk-averse households seek to invest a lower fraction of their wealth in risky assets. As a result, they are more likely to participate under a lower minimum than under a higher one.

Our treatment exposure variable, $Middle_i$, may be subject to additive measurement error due to self-reporting. As we show formally in Appendix B, such measurement error tends to bias the estimate toward zero (i.e., attenuation bias). The exception is if new participants underreport their wealth relative to existing participants. We mitigate this concern by remeasuring $Middle_i$ in two ways. First, we redefine the middle class exclusively as the second quintile of the U.S. wealth distribution and omit households from the third quintile from the sample. Under this definition, upper-class households would need to underreport liquid assets by at least \$36,000 to be misclassified as middle-class. Second, we exclude households whose liquid assets are within a 10% buffer of the boundary between the third and fourth quintiles. This approach removes all cases of mismeasurement less than \$8,400 ($2 \times 0.1 \times 42,000$). Columns (6) and (7) show that the estimates based on these alternative measures of $Middle_i$ equal 0.15 and 0.16. This range lies close to our baseline estimate of 0.14, suggesting that it is not biased by measurement error.

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

4.4 Magnitude of Effect

We use the estimates in Table 2 to decompose the observed growth rate in the total number of robo participants into the component due to the reduction versus that due to other forces. Let g denote the observed growth rate, which we calculate directly from the data. Consider a counterfactual without the reduction, in which middle-class households do not experience a relaxation of minimum-account constraints and, thus, $\mu = 0$. Appendix B shows how the growth rate under

this counterfactual equals

$$g^C = \frac{\psi \mathbb{E}[X_i] + \varrho}{1 - (\psi \mathbb{E}[X_i] + \varrho)}. \tag{4}$$

Our statistic of interest is

$$\eta \equiv g - g^C, \tag{5}$$

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

Table 3 summarizes various calculations of η and of the analogous statistic for growth in middle-class households' robo participation, also derived in Appendix B. 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.[11] The additional estimates from Table 2 imply an increase in the number of middle-class participants between 127% and 148%.

The large effects in Table 3 suggest that middle-class households have significant demand for professional wealth management, but many face minimum-account constraints. In Section 6 , we propose a theory that can quantitatively explain this demand.

5 Robustness

We next assess the internal validity of the baseline results. Specifically, we directly assess the minimum-account constraints channel (5.1), evaluate dynamic confounding channels, such as media attention and advertising (5.2), and discuss various other potential forms of bias (5.3). The results of all these tests support the baseline results' validity.

5.1 Testing the Constraints Channel

We provide regression evidence to confirm the graphical evidence on bunching from Figure 4. In particular, we replace the outcome variable in equation (2) with two indicator variables. The first indicator equals one if the initial deposit is less than \$5,000. The second indicator equals one if the initial deposit equals \$5,000 or is no more than 5% larger.

[11] 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.

Table 4 reports the results. Column (1) shows that new middle-class participants are 29 pps more likely to invest under $5,000 than new upper-class participants. This finding suggests that many new middle-class participants would have liked to invest under $5,000 before the reduction, but the minimum precluded them from doing so. Consistent with this view, the effect is more than twice as strong for middle-class investors from the second quintile, for whom minimum-account constraints are plausibly more severe (column (2)).

Column (3) shows that middle-class households who became participants prior to the reduction were 25 pps more likely to invest right at the minimum than upper-class participants. However, their propensity to do so falls by 32 pps afterward. As before, this effect is stronger for households from the second quintile (column (4)). This finding matches the pre-reduction bunching behavior and its post-reduction dissipation shown in panel (b) of Figure 4. These results again support the idea that middle-class households experience a relaxation of minimum-account constraints.

5.2 Dynamic Confounding Channels

Our data's panel structure allows us to rigorously evaluate whether heterogeneous media attention, targeted advertising, pre-trends across wealth quintiles, or other higher-frequency dynamic effects bias our baseline results. We estimate the following regression equation

$$New\ Participant_{i,t} = \mu\ (Middle_i \times Post_t) + \zeta_i + \varrho_t + u_{i,t}, \tag{6}$$

where i and t index household and week; $Post_t$ indicates if t is greater than the week of the reduction; $New\ Participant_{i,t}$ indicates if i becomes a robo participant in week t, as opposed to the other weeks in our observation window; ζ_i is a household fixed effect; and ϱ_t is a month fixed effect. The parameter μ now equals the effect of the reduction on middle-class households' probability of robo participation in any given week. This interpretation differs from its counterpart in equation (2), where it equals the cumulative effect over the post-reduction period. We first estimate equation (6) as-is and report the results in column (1) of Table 5. Standard errors are two-way clustered by household and week. The reduction increases the weekly probability of becoming a robo participant by 0.7 pps, or, cumulatively, 22 pps over the 32-week post-reduction period (32 $\times$ 0.007). This is on par with the estimated effect in Table 2.

5.2.1 Media Attention and Targeted Advertising

One specific concern is that media attention, advertising, or other changes in visibility around the reduction may disproportionately influence middle-class households, thus biasing our baseline results upward. We address this possibility by collecting additional data on news articles from Google News and on blog posts written by the robo advisor itself. Then, we create two variables: $\textit{Monthly News Articles}_t$, defined as the number of news articles about the advisor published in the month of week t, which proxies for media attention; and $\textit{Monthly Advisor Blogs}_t$, defined as the number of blog posts written by the advisor in the month of week t, which proxies for advertising. We then interact $Middle_i$ with the previous two proxies. This interaction allows middle-class households to respond differently to the dynamic visibility of the robo advisor. We allow for this responsiveness to increase after the reduction (e.g., targeted advertising) by further interacting these proxies with the product $Middle_i \times Post_t$. The corresponding coefficient of interest in columns (2) through (5) of Table 5 is unchanged. This finding suggests that the baseline results do not confound changes in visibility that may disproportionately influence the middle class.

5.2.2 Pre-Trends and Other Dynamic Effects

A more general concern is that our baseline results may confound any dynamic effect that occurs over our observation window and disproportionately affects the middle class. Examples include a secular trend in middle-class households' demand for automated wealth management or changes in industry competition for the middle class. We address this concern by replacing $Post_t$ in equation (6) with a set of indicator variables that equal one if the month of week t is k months before the reduction, denoted $\textit{Months Before}_{t,k}$, or after it, denoted $\textit{Months After}_{t,k}$. The coefficients on the interaction between $Middle_i$ and these indicator variables represent the weekly probability that a middle-class household becomes a robo participant during the indicated month, relative to the reference month of June 2015 (i.e., $\textit{Months Before}_{t,1}$).

The results in column (6) of Table 5 show that the probability of becoming a robo participant increases sharply and significantly for middle-class households exactly in the month of the reduction, July 2015, consistent with Figure 3. By contrast, the middle and upper classes remain on parallel trends over the preceding months, as implied by the insignificant coefficients on the interactions with $\textit{Months Before}_{t,k}$. The precise timing of this increase makes it unlikely that pre-trends in middle-class households' robo participation or other dynamic effects bias the baseline

results. Otherwise, such confounding factors would need to coincide exactly with the month of the reduction, which, per the institutional background in Section 2, is highly unlikely. Column (7) of Table 5 also shows that millennials do not respond to the reduction for reasons apart from wealth, suggesting that our results are not driven by, say, millennials' technological savviness.

5.3 Business Stealing and Gambling

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

6 Life Cycle Model

Introducing a model achieves two purposes that we cannot achieve through a purely reduced-form analysis. First, from a positive perspective, the model explains why the reduction relaxes constraints to the degree documented in Section 4. Absent a model, it is not obvious why households would respond at all, given that they can, in principle, replicate the robo portfolio through self-management. Second, from a welfare perspective, the model allows us to assess the distributional effects of the reduction and, in particular, to study these effects under counterfactual designs of the robo market. We describe the model's setup (6.1) and calibration (6.2) in this section. Positive and welfare implications are in Sections 7 and 8, respectively.

6.1 Setup

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

6.1.1 Preferences

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

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

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

6.1.2 Financial Assets

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

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

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

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

6.1.3 Labor Income

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

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

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

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

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

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

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

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

where B is a function of parameters.[12]

6.1.4 Consolidated Problem

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

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

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

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

$$1-\alpha^f_{i,t}-\alpha^{\mathcal{S}}_{i,t}-\alpha^{\mathcal{A}}_{i,t} \geq 0, \tag{14}$$

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

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

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

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

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

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

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

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

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

6.2 Calibration

Table 6 summarizes the model's parameters and their calibrated values. We first discuss the portfolio parameters and asset pricing factors, shown in panels (a)-(b). Appendix C has details.

6.2.1 Portfolio Parameters

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

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

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

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

We specify F as the three Fama and French (1993) factors with two additional bond factors: the excess returns on U.S. and global bonds, based on Barclays' aggregate bond indices. This 5×1 vector includes many of the factors to which the robo advisor claims to give exposure. Therefore, it likely describes the true return structure more accurately than, say, the CAPM, which we nevertheless consider for robustness. Indeed, Appendix Table A5 shows how robo portfolios have greater net exposure to bonds and to value stocks than their self-managed match. We calibrate the mean and covariance matrix of F using the longest available time series over 1960-2017 and report these values in Appendix Table A6. For reference, the mean and volatility of the market factor equal 7.6% and 14.7%, respectively. Note that calibrating the factor moments to their long-run values may understate the reduction's welfare impact, given evidence of structurally lower volatility since 2010 (e.g., Smith and Timmermann (2021)).

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

$$\sigma_{\epsilon,j}^{\mathcal{P}} = \sqrt{w_j^{\mathcal{P}\prime} \hat{\Sigma}_\epsilon w_j^{\mathcal{P}}}, \quad \beta_j^{\mathcal{P}} = w_j^{\mathcal{P}\prime} \hat{\beta}, \tag{19}$$

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

We pause to summarize the self-managed and matched robo portfolios in our data. Table 7 shows how self-managed portfolios are much less diversified than their robo match. In particular, robo portfolios feature a 30 pps higher Sharpe ratio, which, interestingly, holds for both the middle and upper classes. The higher Sharpe ratio partly reflects a 2 pps higher expected return on robo portfolios. By construction, this higher expected return stems from exposure to priced risk. However, robo portfolios contain less total risk because their idiosyncratic volatility is 11 pps lower. The higher idiosyncratic volatility in self-managed portfolios reflects how they concentrate most of their value in individual stocks and actively-managed funds: the median self-managed portfolio allocates only 8% to broad-based index-linked ETFs, versus 100% for robo portfolios. Appendix Table A7 conveys similar patterns based on other factor models.

Lastly, in the third step, we embed the empirical portfolio parameters in the model. Recalling

that age and liquid assets comprise the model's state variables, we project each of the empirical portfolio parameters on the holder's log liquid assets and a fifth-order polynomial in age. Then, we substitute the fitted values into the parameter vector $\left\{\sigma_{\epsilon,i}^{\mathcal{S}}, \sigma_{\epsilon,i}^{\mathcal{A}}, \beta_i^{\mathcal{S}}, \beta_i^{\mathcal{A}}\right\}$ according to the age and liquid assets of household i. We find similar results when simply substituting the sample average.

6.2.2 Other Parameters

We choose preference parameter values of $\gamma = 9$ and $\delta = 0.96$, consistent with the literature. Appendix Table A9 shows how we obtain similar values when structurally estimating these parameters. We follow Cocco, Gomes and Maenhout (2005) in our calibration of labor income parameters. Accordingly, the deterministic component of income, f_t, is a third-order polynomial in household age, and the coefficients equal those estimated by Cocco, Gomes and Maenhout (2005) for their baseline analysis. Similarly, we parameterize $\sigma_\nu = 0.103$, $\sigma_{\tilde{\xi}} = 0.271$, and $\beta^Y = 0.001$. We parameterize $\phi = 0.001$ to match the share of households in the 2016 SCF earning less than \$10 in total income. The remaining parameter values are: $R^f = 0.2\%$, corresponding to the average one-month Treasury yield over 2010-2017; $t_0 = 25$, $\underline{T} = 65$, and $\overline{T} = 100$, all of which are standard; and p_t, which we calculate using the Center for Disease Control's mortality tables (Xu et al. (2020)).

7 Positive Implications

We first examine whether the model can quantitatively explain the quasi-experimental evidence in Section 4. We calculate how a household of age t with liquid assets $W_{i,t}$ optimally invests under the previous minimum of \$5,000 and, again, under the reduced minimum of \$500. Then, using this change in household-level investment, we compare the reduction's theoretical effect on the robo market with its empirical effect.[14] By construction, the theoretical effect works through a relaxation of minimum-account constraints, just like the empirical effect (e.g., Section 5.1).

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

[14]The theoretical effect on market-level outcomes comes from aggregating household-level policy functions across the bins of age and liquid assets that define the state space, weighting by the share of households in the 2016 SCF within each bin. We do not simulate household investment over the life cycle because we study how a particular quasi-experiment affects the cross-section of households at a given point in time.

Participants) and for those who optimally invest under the reduced minimum (New Participants). This theoretical democratization matches its empirical analogue relatively well, despite the fact that the model essentially has only two free parameters, γ and δ. Appendix Figure A3 shows that this quality of fit is robust to the choice of pricing factors, F.

Table 8 reproduces three other sets of statistics. First, panel (a) reproduces the growth in the number of robo participants reported in Table 3 (i.e., η). The model matches the overall growth rate quite well, though understating the growth in middle-class participation. On the intensive margin, the model predicts the robo portfolio share of new middle and upper-class participants (i.e., $\alpha_{i,t}^{\mathcal{A}}$) remarkably well, as shown in panel (b). This finding is unexpected because we do not target robo share in our calibration. Lastly, panel (c) replicates Figure 5 in terms of wealth classes, as opposed to wealth quintiles.

Collectively, the model quantitatively matches multiple features of the data fairly well. This suggests that the reduction's large empirical effect on middle-class robo participation reflects an optimal response by households with constrained demand for wealth management. Notably, if robo and self-managed portfolios have the same return, $R_{i,t}^{\mathcal{A}} = R_{i,t}^{\mathcal{S}}$, then the model would predict that the reduction has no effect. Therefore, middle-class demand for wealth management must fundamentally stem from differences in diversification (i.e., $\sigma_{\epsilon,i}^{\mathcal{P}}$) or priced risk (i.e., $\beta_i^{\mathcal{P}}$). We decompose the contribution of these two channels in the next section. For now, we conclude that households act optimally conditional on facing these portfolio differences, and we do not need to weigh in on the precise behavioral channel from which these differences stem.

8 Welfare Implications

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

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

as shown in Appendix C.

Like other papers that use life cycle models to study distributional effects (e.g., Gete and Zec-

chetto (2018)), we recalculate equation (20) separately for subpopulations defined by wealth and age. Then, we decompose equation (20) into three terms that reflect the particular gain from changes in diversification, priced risk, and risky share. Explicitly,

$$q_i = \underbrace{\left[\left(\frac{\overline{V}_i\big|_{\alpha,\sigma^{\mathcal{P}}_{\epsilon,i}}}{\underline{V}_i}\right)^{\frac{1}{1-\gamma}} - 1\right]}_{\text{Priced Risk}} + \underbrace{\left[\left(\frac{\overline{V}_i\big|_{\alpha}}{\underline{V}_i}\right)^{\frac{1}{1-\gamma}} - \left(\frac{\overline{V}_i\big|_{\alpha,\sigma^{\mathcal{P}}_{\epsilon,i}}}{\underline{V}_i}\right)^{\frac{1}{1-\gamma}}\right]}_{\text{Diversification}} + \underbrace{\left[\left(\frac{\overline{V}_i}{\underline{V}_i}\right)^{\frac{1}{1-\gamma}} - \left(\frac{\overline{V}_i\big|_{\alpha}}{\underline{V}_i}\right)^{\frac{1}{1-\gamma}}\right]}_{\text{Risky Share}} \tag{21}$$

where $\alpha_i \equiv \alpha_i^{\mathcal{S}} + \alpha_i^{\mathcal{A}}$ is household i's risky share; $\overline{V}_i\big|_{\alpha}$ is i's expected lifetime utility under the minimum of \$500 after constraining risky share to equal its value under the \$5,000 minimum; and, similarly, $\overline{V}_i\big|_{\alpha,\sigma^{\mathcal{P}}_{\epsilon,i}}$ is i's expected lifetime utility with the additional constraint that self-managed and robo portfolios have the same idiosyncratic volatility.

Moving from left to right, the first term in equation (21) equals the welfare gain under a counterfactual in which households cannot increase their risky share and self-managed and robo portfolios only differ in their quantity of priced risk (i.e., $\beta_i^{\mathcal{A}} \neq \beta_i^{\mathcal{S}}$). The second term equals the marginal gain when the two portfolios also differ in idiosyncratic risk (i.e., $\sigma^{\mathcal{A}}_{\epsilon,i} \neq \sigma^{\mathcal{S}}_{\epsilon,i}$), but households still cannot increase their risky share. Notably, these first two terms equal zero for households who do not participate in the stock market before the reduction (i.e., $\alpha_i = 0$). The third term reflects their welfare gain by allowing risky share to increase.

8.1 Distributional Effects by Wealth

Panel (a) of Table 9 reports the average welfare gain for middle and upper-class households who become robo participants after the reduction. The average new middle-class participant gains 2% in lifetime consumption, compared to almost nothing for the average new upper-class participant.[15] For reference, the workhorse models referenced earlier generally consider a one percent gain in lifetime consumption economically significant.

To place a 2% gain in perspective, we calculate the lifetime consumption gain from a "comparison shock" in which the equity premium permanently rises 4 pps (from 7.6% to 11.6%) but the minimum remains fixed at \$5,000. This shock directly benefits households who already par-

[15]The small 0.01% gain reflects the rare case of upper-class households with a very low unconstrained-optimal risky share. These households do not own risky assets before the reduction because their self-managed portfolio contains too much idiosyncratic risk, and the robo portfolio's \$5,000 minimum requires an investment that exceeds their very small unconstrained-optimum. Consequently, the reduction brings these households into the stock market.

ticipate in the stock market, based on the average market beta of 0.93 in self-managed portfolios shown in Appendix Table A5. Indirectly, it also raises their optimal risky share. Together, these channels raise lifetime consumption by 1.7% for the same middle-class households summarized in panel (a). Thus, new middle-class robo participants value the increased accessibility of robo advice by roughly as much they would value a 4 pps higher equity premium without such accessibility. For further reference, their lifetime consumption rises by 1.4% following a complementary comparison shock that increases log labor income by one standard deviation.

8.1.1 Decomposition of Channels

Panel (b) of Table 9 decomposes the total welfare gain according to equation (21). We find that 0.3 pps (15%) reflects an improvement in priced risk exposure, 1.3 pps (65%) reflects better diversification, and 0.4 pps (20%) reflects a higher risky share. The 0.3 pps gain from exposure to priced risk matches empirical evidence that fund managers are compensated for providing access to such risk (Hitzemann, Sokolinski and Tai (2021)). The comparatively small magnitude of this gain reflects how self-managed portfolios and robo portfolios have similar risk exposure, in that the former's expected return lies only 2.2 pps (22%) below than that of the latter, per Table 7. By contrast, self-managed portfolios have over three times as much idiosyncratic volatility as robo portfolios, also per Table 7. Hence, the reduction leads to a larger 1.3 pps welfare gain through diversification.

Lastly, the 0.4 pps gain from a higher risky share principally reflects the gain from becoming a stock market participant. Intuitively, households seek professional management because their self-managed portfolio is underdiversified and, thus, too risky for its expected return. However, accessing professional management requires a risky share of at least $M/W_{i,t}$, which exceeds the unconstrained-optimum for households with modest wealth. Thus, these households face the choice between an underdiversified, self-managed portfolio with a reasonable risky share and a well-diversified, professionally-managed portfolio with an excessive risky share. As a result, many modestly wealthy households may simply prefer not to participate in the stock market. The reduction benefits such households by allowing them to access wealth management and, thus, the stock market, with a less excessive risky share.

8.1.2 Model Extensions

Panel (c) summarizes welfare gains under three model extensions, the details of which are in

Appendix C. First, we introduce a per-period cost of holding the self-managed portfolio equal to $100, or around 10% of the inflation-adjusted cost in Vissing-Jørgensen (2003). Life cycle models typically choose a higher value for this parameter with the intent of capturing the effects of account minimums and underdiversification. We explicitly account for these effects, and so we choose a smaller value that, say, captures time costs associated with rebalancing on one's own.[16] Under this extension, we find that middle-class households experience a larger 2.5% welfare gain, suggesting that the model's parsimony leads to conservative results.

Next, we allow households to borrow at the average interest rate on credit card debt in 2015. Relaxing borrowing constraints in this manner expands households' choice set and so raises their lifetime utility, regardless of the account minimum. Therefore, in relative terms, the reduction should increase welfare by less under this extension because households' pre-reduction utility is higher. Indeed, we find a smaller increase in lifetime consumption of 1.7%. However, this increase is still substantial, suggesting that the reduction improves welfare even if households finance their robo investments with borrowing.

Lastly, we incorporate a defined contribution plan. Like in Campbell et al. (2001), households must allocate 10% of their annual income to this plan and cannot withdraw funds until retirement. The first feature limits households' investible resources, while the second raises the relative value of portfolios that households can liquidate at any time (e.g., the robo portfolio). Together, these effects make a lower account minimum significantly more valuable, leading to a 3.3% welfare gain for the middle-class, 1.3 pps larger than in the baseline model.

Appendix Table A8 supports the robustness of the baseline results by reproducing them under the following parameterizations: structurally estimated preference parameter values; a high discount factor of $\delta = 0.90$; and a 20% correlation between labor income and financial returns.

8.2 Distributional Effects by Age

Table 11 repeats the core exercise in Table 9 by age. Panel (a) shows that the average welfare gain for middle-class participants increases in age. This finding is surprising given that robo advisors claim to "build our products and services for millennials" (e.g., Hutchins (2020)).

[16] For background, workhorse models typically feature a fixed dollar cost of stock market participation, as summarized by Gomes (2020). Without such a cost, these models generally predict that almost all households participate in risky asset markets. In reality, however, only 37% participate (e.g., Appendix Table A3). Our model features a 44% stock market participation rate, which is close to the data. The reason why our model can deliver realistic, limited stock market participation without additional cost parameters is because it incorporates imperfectly diversified portfolios and an account minimum.

We explain this result by first noting that middle-class households' optimal robo participation rate increases in age (Appendix Figure A4). This finding reflects how households accumulate liquid assets as they age and, thus, can invest at the minimum with a lower risky share. Consequently, prior to the reduction, robo non-participants under age 36 have a 77% cumulative probability of eventually participating by retirement, as shown in panel (b) of Table 11. By contrast, households over age 55 who have not yet participated with the robo advisor have only a 17% cumulative probability of eventually participating.

Next, we note that the reduction uniformly increases middle-class households' optimal robo participation rate across the age distribution (Appendix Figure A4). The uniformity of this shift matches our empirical evidence in column (4) of Table 2. The effect accumulates over the life cycle such that, under the reduced minimum, households below age 56 will participate with a probability of 100% by retirement. However, in relative terms, households above age 55 benefit the most because their probability of eventually participating was the lowest before the reduction. In particular, the reduction raises the eventual probability of participation by 74 pps for households over age 55 (i.e., $74 = 91 - 17$), as also shown in panel (b) of Table 11. However, it does so by only 23 pps for those under age 36 (i.e., $23 = 100 - 77$). Thus, the reduction relaxes a "temporary" constraint on the younger middle class but a "permanent" constraint on the older middle class. For this reason, the older middle class experiences a greater welfare gain.

9 Conclusion

We draw two conclusions. First, from a policy perspective, our results exemplify how private wealth management can improve the financial condition of modestly wealthy households. This conclusion comes from studying a large and unexpected reduction in account minimum by a major U.S. automated wealth manager, or robo advisor. The reduction increases the number of robo participants from the middle segments of the U.S. wealth distribution by 110%. This finding suggests that automated wealth management may substitute for government programs that, with mixed rates of success, have attempted to expand the investment opportunities available to the modestly wealthy (e.g., myRA, OregonSaves, NEST).

Second, from the perspective of economic theory, our results support models of bounded rationality in which households act optimally given limits on their ability to invest efficiently on their own. We arrive at this conclusion by quantitatively explaining the previous quasi-experiment

with a life cycle model calibrated to match portfolio-level data. Households optimally seek professional management because they cannot diversify away uncompensated risk as well as a professional manager can. By reducing its minimum, the automated wealth manager enables modestly wealthy households to benefit from professional management, thus improving their welfare by the same amount as would a 4 pps higher equity premium. We leave open the question of why households invest inefficiently on their own, as well as whether these results extend to hybrid wealth management that does not rely fully on automation.

References

Bach, L., Calvet, L. E. and Sodini, P.: 2020, Rich Pickings? Risk, Return, and Skill in Household Wealth, *American Economic Review* . 4

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* . 1

Bachas, P., Gertler, P., Higgins, S. and Seira, E.: 2020, How Debit Cards Enable the Poor to Save More, *The Journal of Finance* . 1

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

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* . 1

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

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

Bombardini, M. and Trebbi, F.: 2012, Risk Aversion and Expected Utility Theory: An Experiment with Large and Small Stakes, *Journal of the European Economic Association* . 1, B.4

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

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

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

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

Campbell, J. Y., Cocco, J. a. F., Gomes, F. J. and Maenhout, P. J.: 2001, Investing Retirement Wealth: A Life-Cycle Model, *Risk Aspects of Investment-Based Social Security Reform*. 6.1, 8.1.2, 9, C.3.3

Campbell, J. Y., Ramadorai, T. and Ranish, B.: 2019, Do the Rich Get Richer in the Stock Market? Evidence from India, *American Economic Review: Insights* . 4

Campbell, J. Y. and Viceira, L. M.: 2002, *Strategic Asset Allocation: Portfolio Choice for Long-term Investors*, Oxford University Press. 4.3

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

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

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

Cocco, J. F., Gomes, F. J. and Maenhout, P. J.: 2005, Consumption and portfolio choice over the life cycle, *The Review of Financial Studies* . 1, 6.1, 6.1.3, 13, 6.2.2, 6, C.1.2

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

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

D'Acunto, F., Prabhala, N. and Rossi, A. G.: 2019, The Promises and Pitfalls of Robo-Advising, *Review of Financial Studies* . 1, 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, 2.1

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

DeFusco, A. A., Johnson, S. and Mondragon, J.: 2020, Regulating Household Leverage, *Review of Economic Studies* . 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

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

Fagereng, A., Guiso, L., Malacrino, D. and Pistaferri, L.: 2020, Heterogeneity and Persistence in Returns to Wealth, *Econometrica* . 4

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

Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T. and Walther, A.: 2021, Predictably Unequal? The Effects of Machine Learning on Credit Markets, *The Journal of Finance* . 1

Fuster, A., Plosser, M., Schnabl, P. and Vickery, J.: 2019, The Role of Technology in Mortgage Lending, *Review of Financial Studies* . 1

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

Gete, P. and Zecchetto, F.: 2018, Distributional Implications of Government Guarantees in Mortgage Markets, *Review of Financial Studies* . 8

Gomes, F.: 2020, Portfolio Choice Over the Life Cycle: A Survey, *Annual Review of Financial Economics* . 1, 16

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

Guiso, L., Sapienza, P. and Zingales, L.: 2008, Trusting the Stock Market, *the Journal of Finance* . 4.2

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

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

Heredia, L., Bartletta, S., Carrubba, J., Frankle, D., Kurihara, K., Macé, B., Palmisani, E., Pardasani, N., Schulte, T., Sheridan, B. and Xu, Q.: 2020, Global Asset Management 2020: Protect, Adapt, and Innovate, *Boston Consulting Group Technical Report* . 1

Higgins, S.: 2020, Financial Technology Adoption. 1

Hitzemann, S., Sokolinski, S. and Tai, M.: 2021, Paying for beta: Leverage demand and asset management fees, *Working Paper* . 1, 8.1.1

Hong, C. Y., Lu, X. and Pan, J.: 2020, FinTech Adoption and Household Risk-Taking, *NBER Working Paper* . 1

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. 8.2

Kacperczyk, M., Nosal, J. and Stevens, L.: 2019, Investor Sophistication and Capital Income Inequality, *Journal of Monetary Economics* . 4

Kaniel, R. and Parham, R.: 2017, WSJ Category Kings - The Impact of Media Attention on Consumer and Mutual Fund Investment Decisions, *Journal of Financial Economics* . 1, 4.2

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

Lusardi, A., Michaud, P.-C. and Mitchell, O. S.: 2017, Optimal Financial Knowledge and Wealth Inequality, *Journal of Political Economy* . 4

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

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

Mihet, R.: 2020, Financial Technology and the Inequality Gap. 4

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

Newey, W.: 2007, Course Materials for Nonlinear Econometric Analysis, *MIT OpenCourseWare* . C.4

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

Piketty, T.: 2014, *Capital in the 21st Century*, Belknap Press. 1

Pilon, M.: 2011/7/30, How 'Separate Accounts' Can Disappoint Investors, *Wall Street Journal* .

Reeves, R. V.: 2015, The Dangerous Separation of the American Upper Middle Class, *Brookings Report* . 3

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

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

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

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

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

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

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

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

Vissing-Jørgensen, A.: 2003, Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions, *NBER Macroeconomics Annual* . 1, 1, 8.1.2, C.3.1

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

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

Figure 1: Shift in Wealth Distribution of Robo Participants

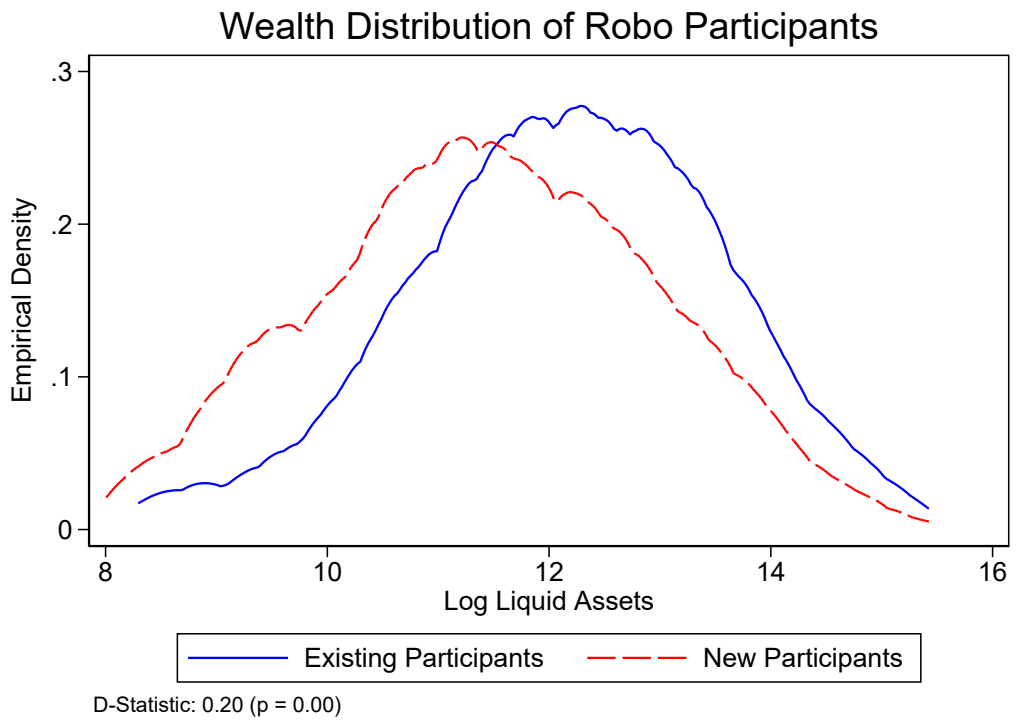

Note: This figure plots the distribution of log liquid assets among households who participated with the robo advisor before the reduction in account minimum (Existing Participants) and who become robo participants after the reduction (New Participants). Liquid assets are defined in Table 1. The distribution is calculated using a kernel density. The D-statistic is based on the Kolmogorov-Smirnov test for equality of distributions.

Figure 2: Change in Representativeness of Robo Wealth Distribution

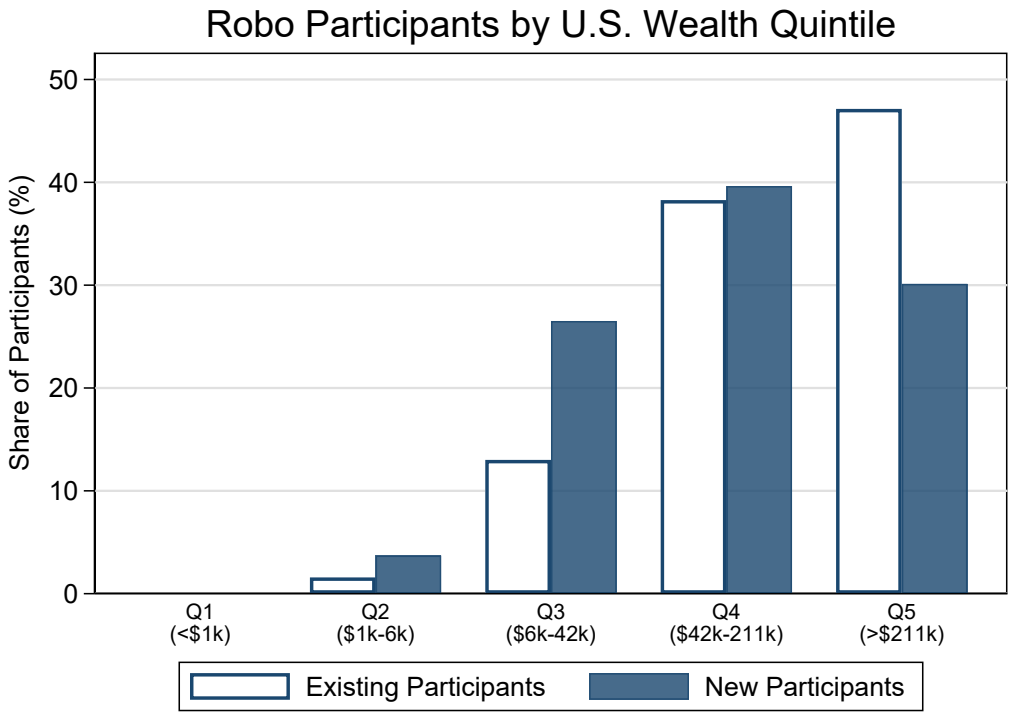

Note: This figure plots the share of robo participants from each quintile of the U.S. wealth distribution. The share is calculated separately for households who participated before the reduction in account minimum (Existing Participants) and who become participants after the reduction (New Participants).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

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

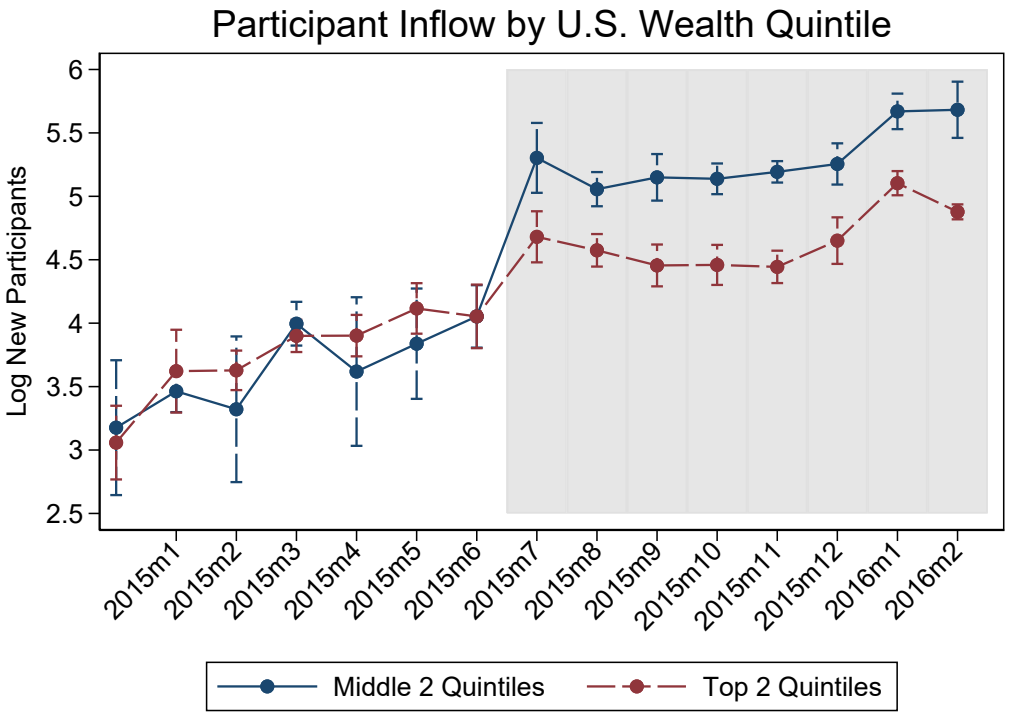

Note: This figure plots the log of the number of new robo participants from the second and third quintiles (Middle 2) and fourth and fifth quintiles (Top 2) of the U.S. wealth distribution, averaged across weeks in each month. The plot is recentered such that the outcome variable is equal across the two wealth classes in June 2015, which allows for an inspection of pre-trends. Wealth consists of liquid assets, defined in Table 1. Wealth quintiles are calculated using the Survey of Consumer Finances (SCF) dataset. The shaded region corresponds to the period after the reduction in account minimum. Brackets are 95% confidence intervals for the monthly average.

Figure 4: Constrained Investment Behavior by the Middle Class

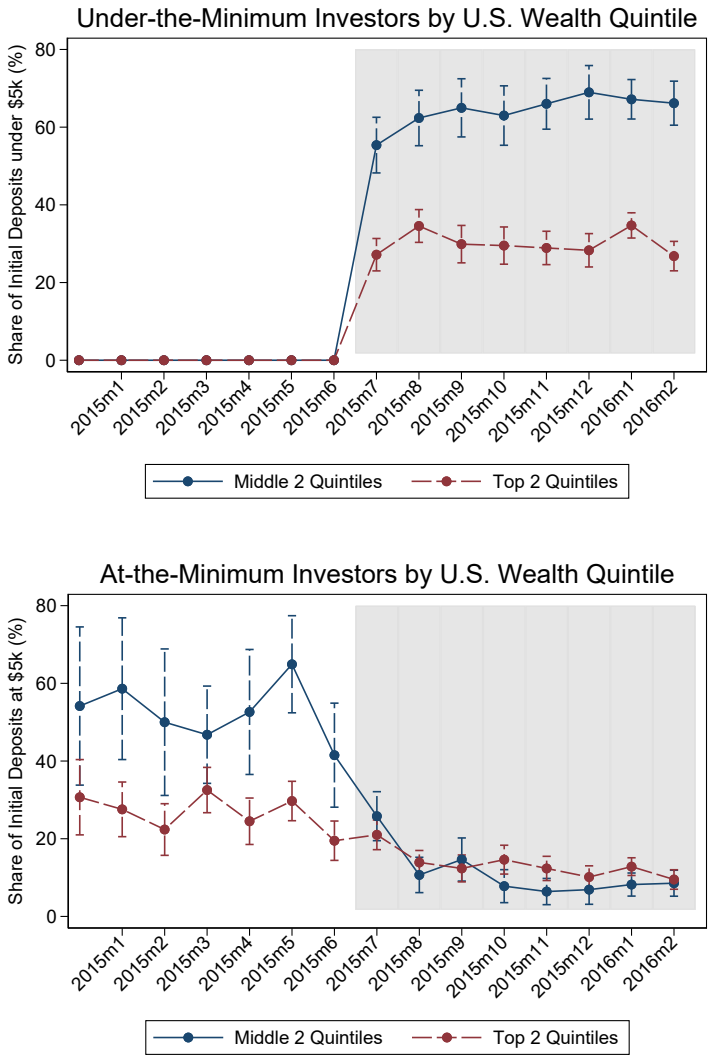

Note: Panel (a) plots the share of new robo participants whose initial deposit is less than the previous account minimum ($5,000) separately for participants from the second and third quintiles (Middle 2) and fourth and fifth quintiles (Top 2) of the U.S. wealth distribution. Panel (b) plots the share whose initial deposit equals the previous account minimum or is no more than 5% higher. Wealth consists of liquid assets, defined in Table 1. Wealth quintiles are calculated using the Survey of Consumer Finances (SCF) dataset. The shaded region corresponds to the period after the reduction in account minimum. Brackets are 95% confidence intervals for the monthly share of participants.

Figure 5: Theoretical Change in Representativeness of Robo Wealth Distribution

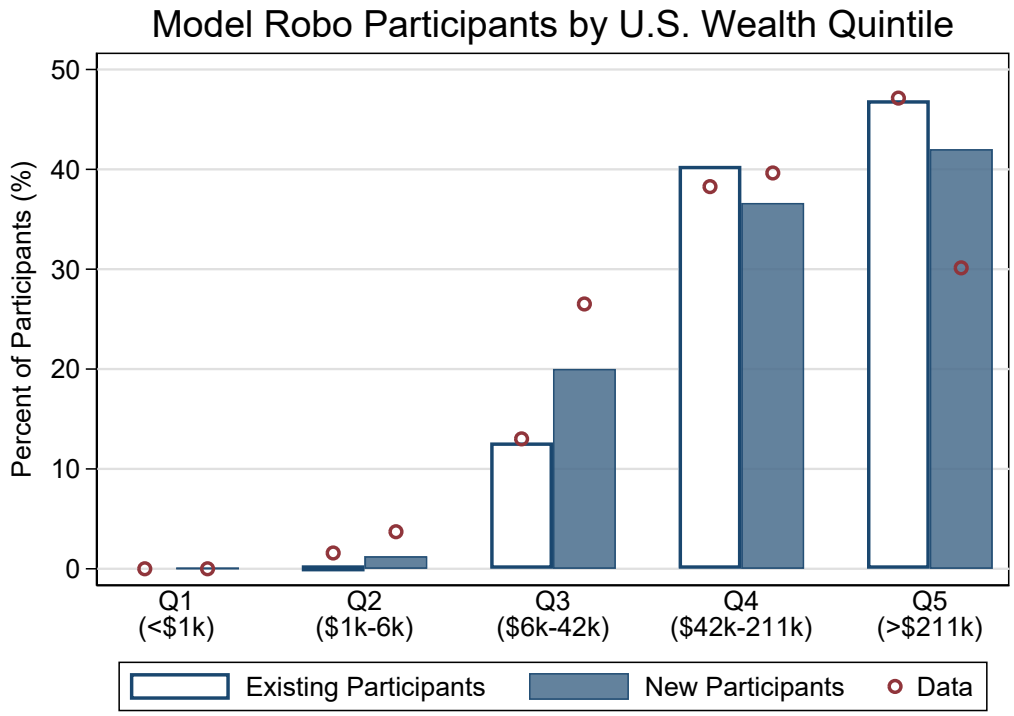

Note: This figure plots the share of robo participants from each quintile of the U.S. wealth distribution, based on the life cycle model in Section 6. Explicitly, figure shows the distribution across wealth quintiles for participants who find it optimal to participate under the previous minimum (Existing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

Table 1: Summary of Robo Participants

	Existing Participants ($N = 4,366$)			New Participants ($N = 5,336$)			
	Mean	Standard Deviation	Median	Mean	Standard Deviation	Median	Difference in Mean
(a) All Households:							
Liquid Assets$_i$ ('000)	436.44	660.82	200	265.21	480.25	100	-171.22
							(0.000)
Income$_i$ ('000)	157.36	110.67	130	116.17	95.9	90	-41.18
							(0.000)
Initial Deposit$_i$ ('000)	33.68	94.54	10	22.56	72.61	5	-11.12
							(0.041)
Age$_i$	35.79	8.72	34	35.4	9.97	33	-0.39
							(0.000)
Middle$_i$	0.15	0.35	0	0.3	0.46	0	0.156
							(0.000)
No Account Closure$_i$	0.95	0.23	1	0.98	0.15	1	0.031
							(0.000)
Subsequent Inflow$_i$	0.9	0.3	1	0.71	0.45	1	-0.185
							(0.000)
(b) Middle Class:							
Liquid Assets$_i$ ('000)	23.23	11.68	25	19.71	11.36	18	-3.527
							(0.000)
Income$_i$ ('000)	92.86	62.21	80	67.14	42.52	60	-25.720
							(0.000)
Initial Deposit$_i$ ('000)	7.6	5.34	5	4.95	12.58	2	-2.652
							(0.000)
Age$_i$	30.33	6.33	29	30.04	7.07	28	-0.293
							(0.339)
No Account Closure$_i$	0.92	0.27	1	0.97	0.18	1	0.043
							(0.000)
Subsequent Inflow$_i$	0.86	0.34	1	0.72	0.45	1	-0.149
							(0.000)

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

Table 2: Democratization of the Robo Market after the Reduction

$Y_i =$	$New\ Participant_i$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Middle_i$	0.219	0.151	0.137	0.180	0.126	0.145	0.155
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$Middle_i \times Age_i$				-0.001			
				(0.384)			
$Middle_i \times Risk\ Averse_i$					0.060		
					(0.041)		
Measure of Middle	Second or Third Quintile					Second Quintile	Middle with Buffer
Controls	No	Yes	Yes	Yes	Yes	Yes	Yes
State FE	No	No	Yes	Yes	Yes	Yes	Yes
R-squared	0.033	0.067	0.097	0.097	0.097	0.078	0.098
Number of Observations	9,349	9,349	9,349	9,349	9,349	7,530	8,982

Note: P-values are in parentheses. This table estimates equation (2), which assesses whether the reduction in account minimum brings middle-class households into the market for automated asset management. Subscript i indexes household. The regression equation is of the form

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

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 (6)-(7) assess the scope for measurement error by remeasuring $Middle_i$ using alternative measures: an indicator for whether i belongs to the second U.S. wealth quintile, after assigning a missing to households from the third U.S. wealth quintile (Second Wealth Quintile); an indicator for whether i belongs to the second or third U.S. wealth quintile, after assigning a missing value to households whose liquid assets are within a 10% buffer of the third quintile (Middle with Buffer). The sample consists all robo participants in the Deposits Dataset. Wealth consists of liquid assets, defined in Table 1. Wealth quintiles are calculated using the Survey of Consumer Finances (SCF) dataset. Household controls are: the log of annual household income; the householder's age; and an indicator for whether the household chooses a lower risk tolerance score than that recommended by the robo advisor, denoted $Risk\ Averse_i$, defined in Table 1. Standard errors are clustered by household.

Table 3: Magnitude of Effect on Robo Participation

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

Note: This table summarizes the observed and counterfactual growth rates in the number of robo participants around the reduction, which assesses the magnitude of the results in Table 2. Column (1) summarizes the observed growth rate in the total number of robo participants, denoted g, and column (2) summarizes the counterfactual growth rate in the absence of the reduction, denoted g^C and defined in equation (4). 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 4: Minimum-Account Constraints as the Mechanism

$Y_i =$	*Under Minimum*$_i$		*At Minimum*$_i$	
	(1)	(2)	(3)	(4)
Middle$_i$	0.294		0.253	
	(0.000)		(0.000)	
Second Quintile$_i$		0.555		0.309
		(0.000)		(0.006)
Third Quintile$_i$		0.269		0.248
		(0.000)		(0.000)
Middle$_i \times$ *New Participant*$_i$			-0.316	
			(0.000)	
Second Quintile$_i \times$ *New Participant*$_i$				-0.467
				(0.000)
Third Quintile$_i \times$ *New Participant*$_i$				-0.302
				(0.000)
New Participant$_i$			-0.149	-0.149
			(0.000)	(0.000)
Controls	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
R-squared	0.156	0.165	0.096	0.097
Number of Observations	5,088	5,088	6,890	6,890

Note: P-values are in parentheses. This table estimates variants of equation (2), which assess the robustness of interpreting households from the second or third U.S. wealth quintiles as constrained by the previous account minimum. Subscript i indexes household. The regression equation is of the form

$$Y_i = \lambda_0 Middle_i + \lambda_1 X_i + \lambda_2 + v_i,$$

where Y_i is a measure of constraints imposed on i by the previous account minimum. We consider measures based on the household's investment behavior: *Under Minimum*$_i$ indicates if i's initial deposit is less than the previous account minimum ($5k); and *At Minimum*$_i$ indicates if i's initial deposit equals the previous account minimum or is no more than 5% higher. Columns (3) and (4) test for a change in bunching behavior by middle-class participants by including the interaction between *Middle*$_i$ and *New Participant*$_i$. The sample in columns (1)-(2) consists of households who become robo participants after the reduction. The sample in columns (3)-(4) expands the sample in columns (1)-(2) to include all households who make their initial deposit over our observation window. The remaining notes are the same as in Table 2.

Table 5: Robustness to Media Attention, Advertising, and Other Dynamic Effects

$Y_{i,t} =$	$New\ Participant_{i,t}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Middle_i \times Post_t$	0.007	0.007	0.008	0.007	0.008		
	(0.000)	(0.000)	(0.000)	(0.000)	(0.002)		
$Middle_i \times Monthly\ News\ Articles_t$		-0.000	-0.000				
		(0.186)	(0.333)				
$Middle_i \times Monthly\ News\ Articles_t \times Post_t$			-0.000				
			(0.759)				
$Middle_i \times Monthly\ Advisor\ Blogs_t$				-0.000	-0.000		
				(0.101)	(0.137)		
$Middle_i \times Monthly\ Advisor\ Blogs_t \times Post_t$					-0.000		
					(0.674)		
$Middle_i \times Months\ Before_{t,3+}$						0.000	0.000
						(0.920)	(0.920)
$Middle_i \times Months\ Before_{t,2}$						-0.002	-0.002
						(0.279)	(0.279)
$Middle_i \times Months\ After_{t,0}$						0.007	0.007
						(0.000)	(0.000)
$Middle_i \times Months\ After_{t,1}$						0.005	0.005
						(0.015)	(0.014)
$Middle_i \times Months\ After_{t,2+}$						0.008	0.008
						(0.000)	(0.000)
$Millennial_i \times Post_t$							-0.000
							(0.689)
Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.011	0.011	0.011	0.011	0.011	0.011	0.011
Number of Observations	504,504	504,504	504,504	504,504	504,504	504,504	504,504

Note: P-values are in parentheses. This table estimates equation (6), which assesses the robustness of the baseline results to a dynamic specification that accounts for various time-varying factors that may disproportionately affect middle-class households. Subscripts i and t index household and week. The regression equation in columns 1-3 is of the form

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

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. Columns (2)-(5) include the interaction between $Middle_i$ and a measure of the robo advisor's visibility: $Monthly\ News\ Articles_t$ is the number of news articles about the robo advisor published in the month of week t, a proxy for media attention; and $Monthly\ Advisor\ Blogs_t$ is the number of blog posts written by the robo advisor in the month of week t, a proxy for advertising. Columns (6)-(7) replace $Post_t$ with an indicator for whether t is k months before or after the reduction, respectively denoted $Months\ Before_{t,k}$ and $Months\ After_{t,k}$, where the reference group consists of the month before the reduction ($Months\ Before_{t,1}$). Column 5 includes the interaction between $Post_t$ and an indicator for whether i is under 35 years old, denoted $Millennial_i$. Note that the millennial generation is defined by being born between 1981 and 1996, such that a household under 35 years of age at the time of the reduction would be a millennial. The set of news articles used to construct $Monthly\ News\ Articles_t$ are the top 150 articles, sorted by relevance, from a Google News search of the advisor's name ("Wealthfront") among articles published in 2015. Standard errors are two-way clustered by household and week. The remaining notes are the same as in Table 2.

Table 6: Model Parameters

Parameter	Value	Source
(a) Portfolio Parameters:		
Idiosyncratic Volatility ($\sigma_{\epsilon,i}$)	Table 7	Portfolio Dataset
Factor Loadings (β_i)	Appendix Table A5	Portfolio Dataset
(b) Asset Pricing Factors:		
Market Factor, Mean	0.076	CRSP
Market Factor, Volatility	0.147	CRSP
Fama-French Factors	Appendix Table A6	French
Bond Factors	Appendix Table A6	Bloomberg-Barclays
(c) Preferences:		
Coefficient of Relative Risk Aversion (γ)	9	Standard
Discount Factor (δ)	0.96	Standard
(d) Labor Income Parameters:		
Age Profile ($f(x)$)	$0.1682 \cdot x - 0.0323 \cdot x^2/10 + 0.002 \cdot x^3/100$	CGM
Permanent Shock Volatility (σ_v)	0.103	CGM
Temporary Shock Volatility ($\sigma_{\tilde{\xi}}$)	0.271	CGM
Loading on Financial Return (β^Y)	0.001	CGM
Probability of Disaster (ϕ)	0.001	SCF
(e) Other Parameters:		
Risk-Free Rate (R^f)	0.002	French
Pre-Retirement Age Range ($[t_0, \underline{T}]$)	[25, 65]	CDC
Range of Survival Rates ($[p(\underline{T}), p(t_0)]$)	[0.865, 0.999]	CDC

Note: This table summarizes the baseline calibration of the life cycle model in Section 6. Panel (a) notes the location of the table summarizing portfolio parameters. Panel (b) summarizes asset pricing factors, presenting the mean and volatility of the market factor and notes the location of the tables summarizing the other factors. Factor moments are calibrated using the means and covariances evaluated over the longest available time series over 1960-2017. Panel (c) summarizes preference parameters. Panel (d) summarizes parameters of the labor income process. Note that: a loading of log labor income on financial returns of $\beta^Y = 0.001$ corresponds to a correlation coefficient of 1%; the probability of labor income disaster is calculated as the share of households in the 2016 SCF with total income less than \$10. Panel (e) summarizes other parameters: the risk-free rate corresponds to the average one-month Treasury yield in 2016; households begin their problem at age t_0, retire after age $\underline{T}$, and leave the model at $\overline{T} = 100$; and the survival rate corresponds to the probability that a household of age t survives until age $t+1$, and it is monotonically decreasing in age. Column (3) reports the source of each value: CRSP denotes the annually-updated stock file from CRSP; French denotes Ken French's website; Bloomberg-Barclays denotes the Bloomberg-Barclays aggregate U.S. and unhedged global bond indices; CGM denotes Cocco, Gomes and Maenhout (2005); CDC denotes the Center for Disease Control's mortality tables; SCF denotes the 2016 Survey of Consumer Finances; and Portfolio Dataset denotes the paper's portfolio dataset summarized in Table 7. Appendix A has details on these data sources.

Table 7: Summary of Self-Managed and Robo Portfolios

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

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

Table 8: Summary of Model Fit

	Model	Data	Source
	(1)	(2)	(3)
(a) Growth in Number of Participants:			
All Participants	11.9%	13.4%	Table 3
Middle-Class Participants	87.1%	107.8%	
(b) Robo Share:			
New Middle-Class Participants	32.7%	30.6%	Deposits Dataset
New Upper-Class Participants	16.8%	16.2%	
(c) Distribution of Robo Participants:			
Pre-Reduction Share from Middle-Class	12.7%	13.4%	Deposits Dataset
Post-Reduction Share from Middle-Class	21.3%	28.4%	

Note: This table summarizes the ability of the life cycle model from Section 6 to fit the data. Panel (a) summarizes the growth in the number of overall and middle-class robo participants, calculated as follows: in the model, we compute the percent increase in the number of participants who find it optimal to participate under the reduced minimum relative to the corresponding number who find it optimal to participate under the previous minimum, separately for all participants and those from the middle class; in the data, we calculate growth rates using the estimates from Table 2 as in Section 4.4. Panel (b) summarizes the average portfolio share allocated to the robo advisor for new middle and upper-class robo participants, calculated as follows: in the model, we compute the average robo portfolio share $\alpha_{i,t}^{A}$ among middle and upper class participants who find it optimal to participate under the reduced minimum; in the data, we compute the ratio of robo investment to liquid assets among middle and upper class households who become participants after the reduction. Panel (c) 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. All model-implied statistics aggregate household-level policy functions across the bins of age and liquid assets that define the state space, weighting by the share of households in the 2016 SCF within each bin. The remaining notes are the same as in Table 2.

Table 9: Welfare Implications of the Reduction

	Increase in Lifetime Consumption		
	All Households	Middle Class	Upper Class
	(1)	(2)	(3)
(a) Baseline Model			
Total Gain	1.93%	1.96%	0.01%
(b) Decomposition			
Priced Risk	0.30%	0.30%	0.00%
Diversification	1.26%	1.29%	0.01%
Risky Share	0.38%	0.37%	0.01%
(c) Model Extensions			
Participation Costs	2.33%	2.51%	0.01%
Borrowing	1.70%	1.73%	0.01%
Defined Contribution Plan	2.93%	3.25%	0.24%
(d) Effect of Comparison Shocks			
+4 pps in Equity Premium	1.76%	1.74%	2.88%
+1 sd in Log Labor Income	1.39%	1.41%	0.40%

Note: This table summarizes the average welfare gain for households who participate with the robo advisor under the reduced minimum but not under the previous minimum, based on the life cycle model in Section 6. Welfare gains are measured by the percent increase in annual consumption under the previous minimum that raises a household's expected lifetime utility by the same amount as the reduction, as in equation (20). Panel (a) summarizes the average of this statistic for all new participants in column (1), for new participants from the middle class in column (2), and new participants from the upper class in column (3). Panel (b) decomposes the total welfare gain into three additive channels that respectively capture the welfare gains from: changes in priced risk exposure (Priced Risk); changes in diversification (Diversification); and changes in risky share in response to the previous two changes (Risky Share). Explicitly, the three channels in panel (b) are defined in equation (21). Panel (c) summarizes the average welfare gain under models with the following extensions: a per-period cost of $100 when holding the self-managed portfolio (Participation Costs); the ability to borrow up to 30% of one's liquid assets at the average rate on credit card debt in 2015 of 12% (Borrowing); and, following Campbell et al. (2001), a requirement to allocate 10% of one's income to a defined contribution plan that delivers the same annual return as the risk-free asset and cannot be liquidated until retirement (Defined Contribution). Panel (d) summarizes the welfare gain from alternative shocks that occur under the previous minimum: a permanent 4 pps increase in the expected excess return on the U.S. stock market; and an increase in log labor income equal to one standard deviation of the sum of permanent ($\nu_{i,t}$) and temporary ($\xi_{i,t}$) labor income shocks. Averages are calculated across the bins of age and liquid assets that define the state space, weighting by the share of households in the 2016 SCF within each bin. The remaining notes are the same as in Table 2.

Table 10: Welfare Implications by Degree of Portfolio Personalization

	Moderately Risk Averse		Highly Risk Averse	
	Middle Class	Upper Class	Middle Class	Upper Class
	(1)	(2)	(3)	(4)
(a) Increase in Lifetime Consumption from Reduction:				
Personalization by Age	1.58%	0.01%	0.67%	0.02%
Personalization by Age and Wealth	1.96%	0.01%	1.27%	0.05%
Personalization by Age, Wealth, and Risk Aversion	2.05%	0.01%	1.43%	0.02%
(b) Change in Probability of Robo Participation:	46.50%	1.91%	74.28%	1.93%

Note: This table summarizes the average welfare gain associated with the reduction in account minimum when the robo advisor offers varying degrees of personalization, based on the life cycle model in Section 6. Welfare gains are measured as in Table 9 by the percent increase in annual consumption under the previous minimum that raises a household's expected lifetime utility by the same amount as the reduction, per equation (20). The average is calculated across households who participate with the robo advisor under the reduced minimum but not under the previous minimum. Columns (1)-(2) report this average for middle and upper-class households using a coefficient of relative risk aversion of $\gamma = 9$ (Moderate Risk Aversion), and columns (3)-(4) do similarly using $\gamma = 11$ (High Risk Aversion). Panel (a) summarizes welfare gains when the robo advisor offers one of the following levels of personalization: the allocation varies only by age, where, for each age, the allocation equals the average across wealth for the age-by-wealth pairs imputed from the portfolio dataset in Section 6.2.1; the allocation varies by both age and liquid wealth, where, for each age-by-wealth pair, the allocation equals that imputed from the portfolio dataset, as described in Section 6.2.1 (Robo Portfolio, Unadjusted for Risk Aversion); and the allocation varies by age, liquid wealth, and subjective risk aversion, where, for each age-by-wealth-by-risk-aversion triplet, the allocation equals that imputed from the portfolio dataset as in Section 6.2.1 separately for each value of the empirical measure of high risk aversion (Robo Portfolio, Adjusted for Risk Aversion). The empirical measure of high risk aversion is an indicator for whether the household chooses a lower risk tolerance score than that recommended by the robo advisor, previously denoted $Risk\ Averse_i$ in Table 2. When adjusting for risk aversion, columns (3)-(4) summarize welfare gains based on the robo portfolio calibrated using the subsample of households with high empirical risk aversion, and columns (1)-(2) do analogously using the subsample without high empirical risk aversion. Panel (b) summarizes the model-implied change in the probability of robo participation conditional on eventually participating. The remaining notes are the same as in Table 9.

Table 11: Heterogeneous Welfare Implications by Age

	Age 25-35		Age 36-55		Age 56-65	
	Middle	Upper	Middle	Upper	Middle	Upper
	(1)	(2)	(3)	(4)	(5)	(6)
(a) Increase in Lifetime Consumption from Reduction:	1.81%	0.01%	1.94%	0.01%	2.23%	0.00%
(b) Probability of Becoming Robo Participant by Retirement:						
Pre-Reduction	76.65%		56.10%		17.29%	
Post-Reduction	100.00%		100.00%		91.47%	

Note: This table summarizes the average welfare gain across the age distribution for households who participate with the robo advisor under the reduced minimum but not under the previous minimum, based on the life cycle model in Section 6. Welfare gains are measured as in Table 9 by the percent increase in annual consumption under the previous minimum that raises a household's expected lifetime utility by the same amount as the reduction, per equation (20). Panel (a) summarizes the average of this statistic for middle and upper class households within each of the indicated bins of the age distribution. Panel (b) summarizes the cumulative probability that a robo non-participant of age t becomes a robo participant by age $\underline{T} = 65$. The probability equals $P_t^R = \sum_{\tau=t+1}^{\tau=\underline{T}} \Delta P_t \prod_{j=t+1}^{j=\tau-1}(1-\Delta P_j)$, where P_t is the probability that a household of age t participates with the robo advisor, based on the average across bins of liquid assets in the 2016 SCF weighted by the bin's population share, and $\Delta P_t \equiv P_t - P_{t-1}$. By definition, $P_{\underline{T}}^R = 0$ because a non-participant of age $\underline{T}$ retires in the following year. The remaining notes are the same as in Table 9.

Online Appendix

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

A Data Appendix

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

A.1 Deposits Dataset

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

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

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

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

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

A.2 Portfolio Dataset

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

Specifically, candidate clients would provide their log-in credentials for their outside brokerage account. Then, the robo advisor would take a snapshot of the account holdings and run an advice-generating algorithm on it. This produces a set of snapshots of households' non-robo accounts. While the advice algorithm ran, our data provider would ask the household to answer its standard questionnaire, which is the source of our demographic variables.[17] Finally, at the conclusion of the report, our data provider would tell the household the portfolio she would receive as a client and give her the option to fund a robo portfolio. This produces a matched, counterfactual robo portfolio for each household in the sample. The tool was launched in January 2016, and our sample contains 2,654 snapshots taken between January 2016 and November 2016.

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 7.

A.3 Auxiliary Datasets

A.3.1 Survey of Consumer Finances

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

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

[17] As mentioned in the text, merging the deposits and portfolios datasets is not necessary for our analysis. In principle, we could merge the two dataests on age, income, liquid assets, state of residence, and risk tolerance, but we cannot merge by a unique household identifier for privacy concerns. This exercise results in a relatively low match rate of 3%, which may reflect how the online tool was launched in January 2016, whereas the deposits dataset ends in February 2016. In addition, the low match rate may reflect how the online tool only applied to non-robo portfolios held at one of the following institutions: Charles Schwab, Fidelity, ETRADE, TD Ameritrade, Vanguard, Scottrade, Merrill Lynch, or Morgan Stanley.

A.3.2 Asset Pricing Datasets

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

A.3.3 Other Auxiliary Datasets

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

A.4 Description of Variables

Our empirical analysis relies on the following variables:

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

B Econometric Appendix

We provide the details for aggregate results presented in Table 3 (B.1). Then we explain how the estimation of equation (2) can be affected by measurement error (B.2). Lastly, we discuss the robustness of our main results from Section 4 to the effects of business stealing (B.3) and gambling motives (B.4).

B.1 Aggregation Exercise

We derive the expression for the effect of the reduction on the total number of robo participants shown in Section 4.4. Note that the observe growth rate in the total number of robo participants can be directly calculated from the data as

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

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

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

where $All\ Participants = New\ Participants + Existing\ Participants$ is the sum of new and existing robo participants; and $\mathbb{E}\left[New\ Participant_i\right]$ 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

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

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

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

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

or, equivalently,

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

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

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

which we can rewrite as

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

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

$$\widehat{\mu} = \mu \left(1 - \frac{\text{Var}\left[\varepsilon_i\right] + \mathbb{E}\left[Middle_i \times \epsilon_i\right]}{\text{Var}\left[\widehat{Middle}_i\right]} \right) + \frac{\mathbb{E}\left[u_i \times \varepsilon_i\right]}{\text{Var}\left[\widehat{Middle}_i\right]}. \tag{B7}$$

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 Business Stealing

New middle-class robo participants may have planned to invest with a competitor robo advisor during the post-reduction period, but the reduction prompted them to invest with Wealthfront instead. In this case, our results would reflect business stealing rather than democratization of automated wealth management. We assess this possibility by using data from the SEC's Form ADV to plot new participants at other standalone robo advisors, namely Betterment and Personal Capital. While the Form ADV data have limitations described in Section 3, they are the best source of data for this exercise, short of having microdata from each major U.S. robo advisor. The results in Appendix Figure A2 show very little decline in new participation at Wealthfront's competitors, measured by the log of the change in number of clients, from the pre-reduction to the post-reduction periods. This observation suggests that the reduction indeed expands access to automated wealth management, rather than simply reallocating participants across robo advisors.

B.4 Gambling

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

C Theory Appendix

We describe the solution of the life cycle model in Section 6 (C.1), derive the welfare measure studied in Section 8 (C.2), provide details on the model extensions also studied in Section 8 (C.3), and structurally estimate the model's preference parameters as referenced in Section 6.2 (C.4).

C.1 Model Solution

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

C.1.1 Consolidated Problem

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

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

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

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

Indirect utility has the familiar form

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

with

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

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

C.1.2 Numerical Algorithm

Our numerical algorithm is standard and follows the methods typically used in workhorse life cycle models (e.g., Cocco, Gomes and Maenhout (2005)). First, we solve equation (C1) for age $\underline{T}$ as a function

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

Following standard practice, we optimize by grid search, and so we avoid selecting local optima. Accordingly, we discretize the control variables: $\alpha^{f}_{i,t}, \alpha^{\mathcal{S}}_{i,t}$, and $\alpha^{\mathcal{A}}_{i,t}$. The control variable grid omits choices that violate one of the constraints (8)-(16). As mentioned in the text, we simplify the model's computational complexity by assuming households cannot hold the self-managed and automated portfolios concurrently: $\min\left\{\alpha^{\mathcal{S}}_{i,t}, \alpha^{\mathcal{A}}_{i,t}\right\} = 0$. We obtain similar results without this simplification because it is rarely optimal to hold both at the same time. We also reduce computational complexity by assuming households must maintain a minimum balance of M with the automated asset manager, whereas, in reality, households only need to make an initial deposit of M. Otherwise, we must keep track of $\alpha^{\mathcal{A}}_{i,t}$as an additional state variable because it determines the lower bound on a household's robo investment. Namely, under an initial deposit requirement, households do not need to top-up their balance to M if market forces push their account balance below this threshold. We assess the validity of this simplification by solving the model under the more realistic yet intensive setup with a minimum deposit requirement, finding similar results as in Table 9. Intuitively, the high expected return on the robo portfolio makes cases of a top-up relatively rare.

Next, after solving $V_{\underline{T}}(W_{i,\underline{T}})$, we iterate backward, solving $V_{\underline{T}-1}(W_{i,\underline{T}-1})$ and so forth until we arrive at the initial period, t_0. For each age t, we approximate $V_{t+1}(W_{i,t+1})$ using a cubic 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. This approximation enables the utility function to retain prudence, and it is well-behaved for a suitably fine discretization of the state space. We solve $V_t(W_{i,t})$ using the labor income parameters shown in Table 6, setting income equal to the median income in the 2016 SCF for the baseline cohort studied in Cocco, Gomes and Maenhout (2005).

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

C.2 Welfare Measure

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

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

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

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

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

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

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

C.3 Model Extensions

Section 8.1 considers three model extensions that we now describe. Panel (c) of Table 9 summarizes welfare gains under these extensions.

C.3.1 Participation Costs

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

C.3.2 Borrowing

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

C.3.3 Defined Contribution Plan

The Defined Contribution Plan extension requires households to allocate 10% of their income to a plan that delivers the same annual return as the risk-free asset, R^f. This setup and parameterization follows Campbell et al. (2001). Liquid wealth for household i of age t is defined to include the value of plan assets accumulated to that point. As in Campbell et al. (2001), households cannot withdraw funds from the plan until age $\underline{T} + 1$.

C.4 Structural Parameter Estimation

We assess the validity of the preference parameter values in Table 6 by estimating these parameters through a generalized method of moments estimator (GMM). The 2×1 parameter vector is (γ, δ). The 10×1 moment vector consists of two 5×1 vectors: the share of robo participants from each of the five U.S. wealth quintiles Q under the previous account minimum, which we denote $\underline{\theta}_Q$; and the analogous share under the reduced minimum, which we denote $\overline{\theta}_Q$. Let

$$\Theta = [\underline{\theta}_1, ...\underline{\theta}_5, \overline{\theta}_1, ..., \overline{\theta}_5]'$$

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

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

where Ψ is a 10×10 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 $\hat{\gamma} = 9.1$ and $\hat{\delta} = 0.945$ as shown in Appendix Table A9. We obtain estimates of $\hat{\gamma} = 8.5$ and $\hat{\delta} = 0.990$ when weighting moments equally. We compute standard errors by bootstrapping the entire estimation procedure.

D Asset Pricing Appendix

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

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

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

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

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

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

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

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

Equations (D2)-(D4) are: the standard capital asset pricing model (CAPM), the Fama-French Three Factor Model, and a five-factor model augmenting the Fama-French model with U.S. and global bond returns. Our data on monthly returns come from the Center for Research in Security Prices (CRSP) and Ken French's website, as described in Appendix A. We use the sample mean to calibrate the factor risk prices, multiplied by 12 to obtain an approximate annual value. Analogously, we use the sample covariance matrix to calibrate the covariance of the factor vector. The moments of the factor vectors in equations (D2)-(D4) are shown in Appendix Table A6.

Additional Figures and Tables

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

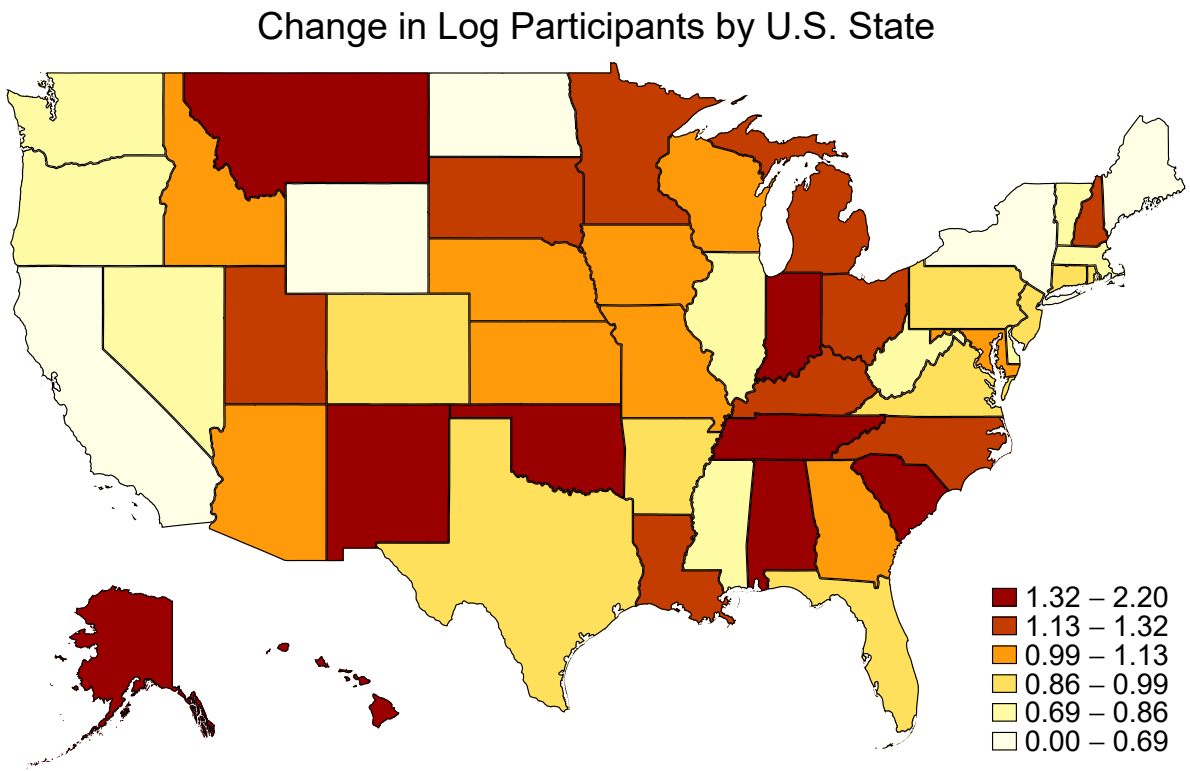

Note: This figure plots the change in the log of the number of robo participants from each U.S. state. The change is from the pre-reduction period (December 1, 2014 to July 7, 2015) to the post-reduction period (July 7, 2015 to February 29, 2016).

Figure A2: Reallocation of Participants across Robo Advisors

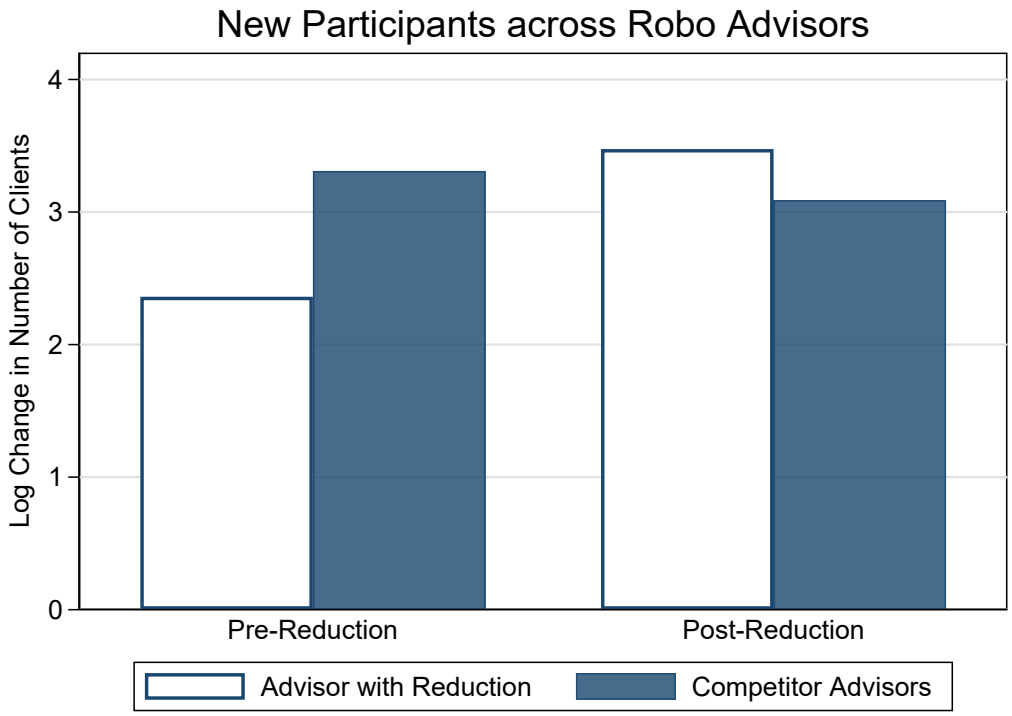

Note: This figure plots the log of the change in the number of clients across robo advisors, in thousands, which assesses whether the reduction increases robo participation or simply reallocates robo participants across advisors. The change is calculated separately for the robo advisor that reduced its account minimum, Wealthfront, and for its competitors combined. The left two columns plot this change over the pre-reduction period (Q4, 2014 to Q2, 2015), and the right two columns plot this change over the post-reduction period (Q2, 2015 to Q1, 2016). 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.

Figure A3: Theoretical Change in Robo Wealth Distribution by Factor Model

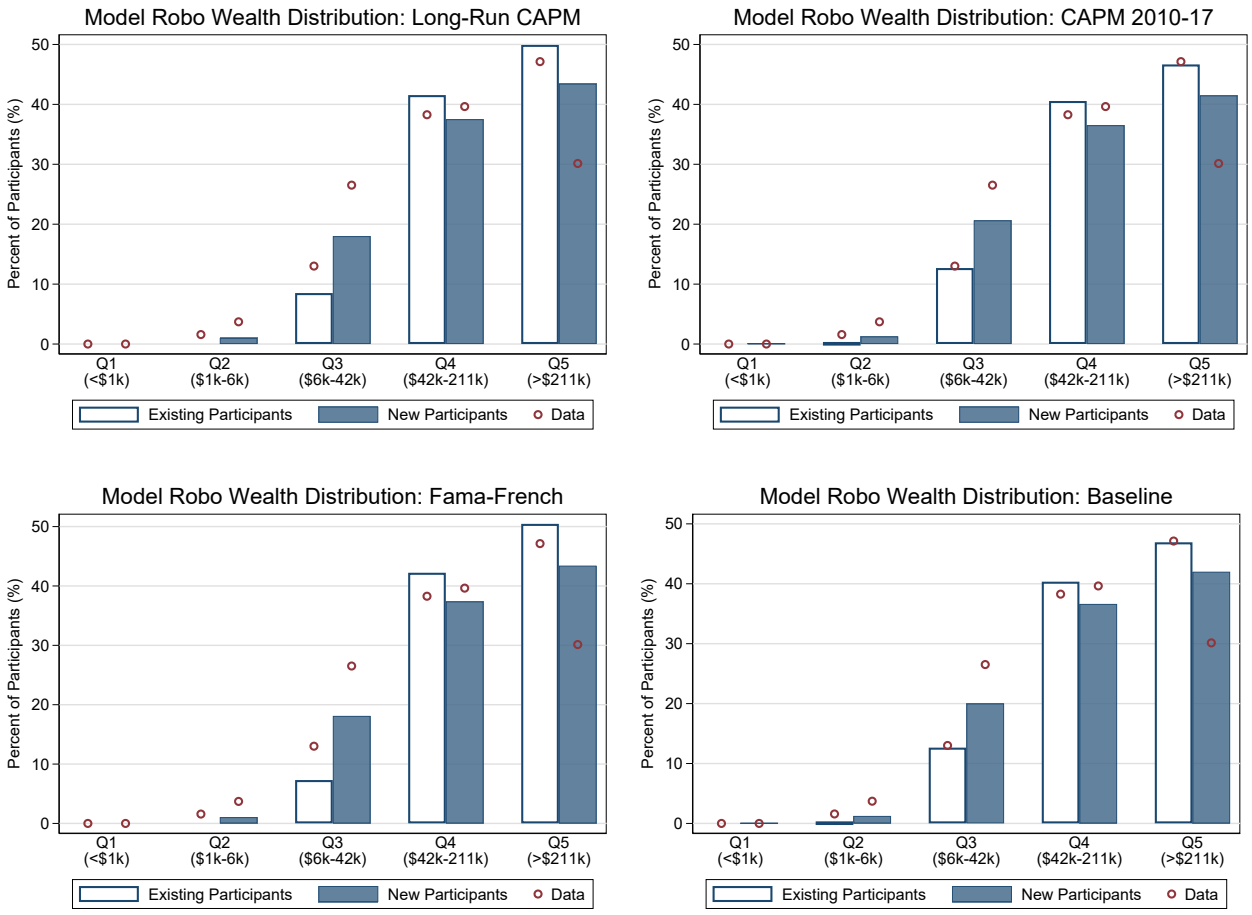

Note: This figure assesses the robustness of Figure 5 by plotting the model-implied share of robo participants from each quintile of the U.S. wealth distribution for various asset pricing factor models. The factor model in the upper two panels is the CAPM, calibrated using the 1960-2017 period in the upper-left panel and the 2010-2017 period in the upper-right panel. The factor model in the lower-left panel is the Fama-French Three Factor Model (Fama-French). The factor model in the lower-right panel is the baseline Fama-French model augmented with U.S. and global bond returns, and so the lower-right panel shows the same figure as in Figure 5. The moments of the pricing factors are summarized in Appendix Table A6. The remaining notes are the same as in Figure 5.

Figure A4: Optimal Middle-Class Robo Participation Rate by Age

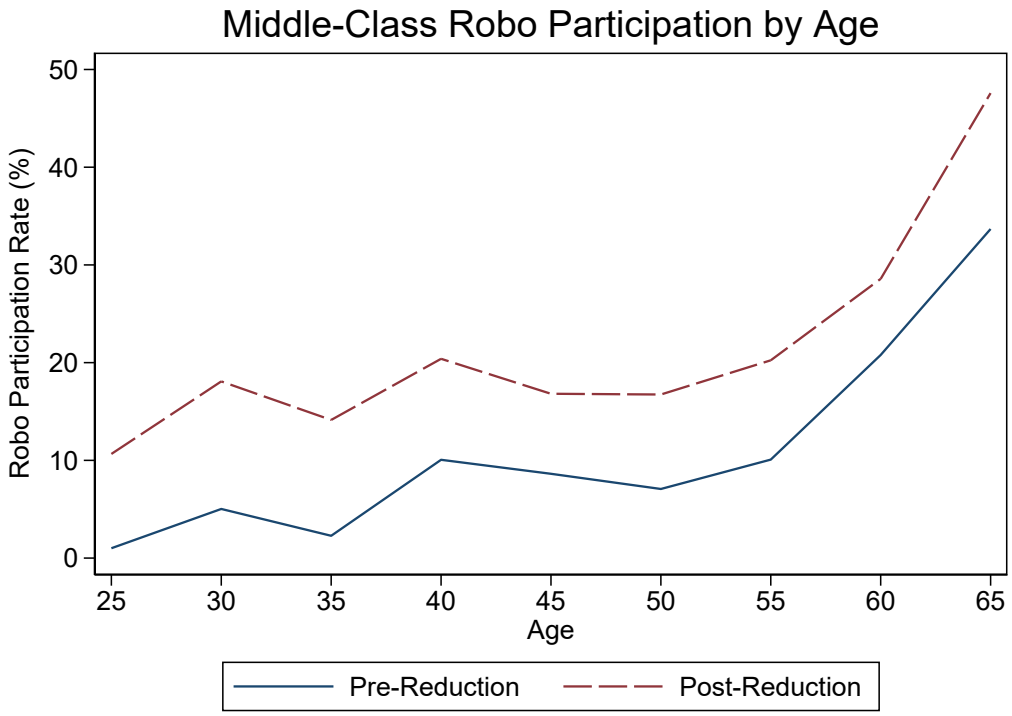

Note: This figure plots the share of middle-class households who find it optimal to participate with the robo advisor by five-year age brackets, based on the life cycle model in Section 6. The blue solid curve plots this share under the previous account minimum ($5,000), and the red dashed curve plots this share under the reduced minimum ($500). The share is calculated by averaging across bins of age and liquid assets that define the state space, weighting by the share of households in the 2016 SCF within each bin. Since older households empirically have higher wealth, this weighting encodes the fact that households accumulate wealth as they age. The remaining notes are the same as in Table 2.

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

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

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

Table A2: Summary of Robo Portfolios

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

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

Table A3: Summary of U.S. Wealth Quintiles

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

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

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

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

Note: P-values are in parentheses. This table reestimates the main specifications in Tables 2 and 4 after clustering standard errors by state of residence.

Table A5: Summary of Non-Robo and Robo Portfolio Loadings

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

Note: P-values are in parentheses. This table summarizes the factor loadings for self-managed and matched robo portfolios, based on the Fama-French Three Factor Model augmented with U.S. and global bond returns. Subscript i indexes portfolio. Each row summarizes the loading on a different factor: Market is the return on the CRSP Value-Weighted Index, net of the risk-free rate; HML is the spread in monthly return between high book-to-market stocks and low book-to-market stocks; SMB is the spread in monthly return between stocks with a small market capitalization and a big market capitalization; GLB is the monthly return on the Barclays Global Aggregate Bond Index Unhedged, net of the risk-free rate; and USB is the monthly return on the Barclays Aggregate U.S. Bond Index Unhedged, net of the risk-free rate. The remaining notes are the same as in Table 7.

Table A6: Covariances and Means of Pricing Factors

Panel (a): Covariance Matrix

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

Panel (b): Means

Market	SMB	HML	GLB	USB
0.076	0.021	0.042	0.060	0.062

Note: This table summarizes the covariance matrix and mean of the baseline asset pricing factor vector, defined in Appendix Table A5. Panel (a) summarizes the covariance matrix. Panel (b) summarizes the mean. Each value in the table is calculated based on the longest available time series over 1960-2017. Over the 2010-2017 period, the volatility of the market factor equals 12.3% and the mean equals 12.6%. The remaining notes are the same as in Appendix Table A5.

Table A7: Summary of Self-Managed and Robo Portfolios by Factor Model

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

Note: P-values are in parentheses. This table assesses the robustness of Table 7 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 7.

Table A8: Robustness of Welfare Implications

	Increase in Lifetime Consumption (%)		
	All Households	Middle Class	Upper Class
	(1)	(2)	(3)
Parameterization			
GMM Estimates ($\gamma = 9.1$, $\delta = 0.94$)	1.58%	1.65%	0.01%
High Impatience ($\delta = 0.90$)	1.95%	1.98%	0.08%
High Labor Income Loading ($\beta^Y = 0.13$)	1.79%	1.91%	0.01%

Note: This assesses the robustness of the welfare implications in panel (a) of Table 8 to alternative parameterizations. Unless otherwise indicated, the values of the remaining parameters are the same as the baseline values in Table 6. Note that a loading of log labor income of financial returns of $\beta^Y = 0.13$ corresponds to a correlation coefficient of 20%. The remaining notes are the same as in Table 8.

Table A9: GMM Estimates of Preference Parameters

Parameter	Estimate			
	(1)	(2)	(3)	(4)
Coefficient of Relative Risk Aversion (γ)	9.1	8.5	8.8	11.2
	[8.868, 9.332]	[8.242, 8.758]	[TBD]	[TBD]
Discount Factor (δ)	0.945	0.990		
	[0.928, 0.962]	[0.970,1.010]		
Sample	Full	Full	Moderate Risk Aversion	High Risk Aversion
Weights	GMM Optimal	Equal	GMM Optimal	GMM Optimal

Note: This table estimates the life cycle model's preference parameters using the generalized method of moments estimator described in Appendix C. The 10×1 moment vector consists of two 5×1 vectors: the share of robo participants from each of the five U.S. wealth quintiles under the previous account minimum; and the analogous share under the reduced minimum. Column (1) estimates the parameters when weighting each empirical moment by its inverse standard error (GMM Optimal). The standard error is replaced by a small positive constant (10^{-4}) for the moments related to the share of participants from the first U.S. wealth quintile, since this share has no empirical variance. Column (2) estimates the parameters when equally weighting each empirical moment. Column (3) restricts the sample used to calculate the empirical moments to households who do not choose a lower risk tolerance score than that recommended by the robo advisor, and column (4) analogously restricts the sample to households who do choose a lower risk tolerance score. Columns (3)-(4) constrain $\delta = 0.96$, as in Table 6. Standard errors are bootstrapped. Brackets correspond to 95% confidence intervals under an asymptotically normal distribution.