

Automation and Inequality in Wealth Management*

Michael Reher[†] and Stanislav Sokolinski[‡]

August 2021

Abstract

We show how automation (i.e., robo advisors) expands access to wealth management and improves middle-class households' welfare. Using novel microdata from a major U.S. robo advisor, we study a quasi-experiment in which the advisor reduces its minimum investment by 90%. The reduction relaxes funding constraints on middle-class households, increasing their participation with the advisor by 110%. We rationalize this behavior with a life cycle model calibrated using portfolio-level data. Middle-class households optimally participate after the reduction because they cannot invest efficiently alone. Their welfare gains equate to a 4 pps higher equity premium and are higher for those near retirement.

Keywords: FinTech, Financial Advice, Portfolio Delegation, Inequality

JEL Classification: G11, G24, D3, O3

*This paper was previously circulated under the title "Does FinTech Democratize Investing?". We thank Christoph Aymanns, John Campbell, Francesco D'Acunto, Ben Friedman, Jakub Jurek, Iva Kalcheva, Jun Liu, Matteo Maggiori, Apurva Mehta, Will Mullins, Alessandro Previtero, Enrichetta Ravina, Jonathan Reuter, Alberto Rossi, Andrei Shleifer, Celine Sun, Ross Valkanov, Boris Vallee, anonymous referees, and seminar participants at the 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 accessible wealth management.² Related work finds that robo advisors benefit affluent investors who already have access to professional management (e.g., [D’Acunto, Prabhala and Rossi 2019](#); [Loos et al. 2020](#); [Rossi and Utkus 2020](#)). We ask whether these gains extend to a much larger and less-wealthy segment of the population.

We show that relaxing funding constraints raises middle-class households’ participation in automated wealth management, thereby improving their welfare through diversified exposure to priced risk. First, using novel microdata and a quasi-experimental research design, we find that reducing minimum investment requirements disproportionately raises participation in wealth management by households from the lower-middle quintiles of the U.S. wealth distribution. Our battery of robustness exercises confirms that the reduction relaxes funding constraints on these households. 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 find that middle-class households optimally respond to the reduction because they struggle to diversify on their own. In particular, the reduction improves middle-class households’ welfare by the same amount as a 4 pps increase in the equity premium, and, surprisingly, those over age 55 gain the most. We conclude that automation, in the form of robo advisors, leads to Pareto improvements by reducing inequality in wealth management.

Empirically, we seek to identify the effect of funding constraints on participation in automated wealth management. Our 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 novel

¹Wealth 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).

²Quoting 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).

microdata directly from a major U.S. robo advisor. Our data come in two parts. In one dataset, we observe the demographic background, investment activity, and liquid assets, inclusive of retirement accounts, of households who participate with the advisor. This information enables our empirical analysis of the distributional effects of automated wealth management. In the second dataset, we observe pairs of portfolios for households interested in becoming robo participants: 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.

To identify the effect, we study 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. Moreover, it generates a clean source of variation with which to test our research hypothesis because it occurs within the same wealth manager, thereby allowing us to hold manager-specific effects fixed. In particular, we avoid bias from a naive comparison between automated and non-automated wealth managers, since robo advisors may attract households who vary in other dimensions that correlate with wealth (e.g., technological savviness).

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 investment constraints due to the reduction. Accordingly, we find that middle-class households are 14 pps more likely to participate with the robo advisor after the reduction, relative to the upper class. We show that this estimate implies a 110% increase in the total number of middle-class robo participants. This finding is robust

³In popular parlance, the term "lower-middle class" typically describes what we call the "middle class", and "upper-middle class" typically describes what we call the "upper class" (Reeves (2015)).

to various measures of wealth and, thus, not biased by measurement error.

To ensure we identify the desired effect, we examine whether a relation of funding constraints drives the results. Consistent with this view, the majority of middle-class households who became robo participants prior to the reduction “bunched” their investment right 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 less than \$5,000. Based on a wide variety of robustness 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 other effects specific to demographics or risk attitude.

The reduction’s large empirical effect raises both positive and normative questions. From a positive perspective, it is not obvious how to reconcile this finding with economic theory. At one extreme, workhorse life cycle models predict that middle-class households can invest efficiently on their own and, thus, would not respond to the reduction (e.g., [Cocco, Gomes and Maenhout \(2005\)](#)). At the other extreme, a behavioral model in which households invest for entertainment would predict no effect, since the robo advisor allows little deviation from its “boring”, canonical allocations (e.g., [Kalda et al. \(2021\)](#), [Ben-David et al. \(2021\)](#), [Welch \(2020\)](#), [Barber et al. \(2020\)](#)). From a normative perspective, wealth managers have historically not catered to the middle class, and so little is known about whether and how they improve middle-class welfare.

We address these questions through the lens of a workhorse life cycle model with two novel ingredients. First, households choose between managing their own portfolio or delegating it to a professional. The two options differ in the quantity of priced versus idiosyncratic risk. We use our rich data on self-managed and counterfactual robo portfolios to realistically calibrate these quantities across the joint distribution of household age and wealth. Based on this calibration, robo portfolios have a 2 pps higher expected return than self-managed ones, which comes from priced risk exposure (e.g., value, bonds). However, robo portfolios have less total risk, reflecting an 11 pps lower idiosyncratic volatility. The model’s second distinguishing ingredient is a minimum investment for professional management.

The model formalizes the following intuition. Per our calibration results, households have limited ability to independently diversify idiosyncratic risk (e.g., [Calvet et al. \(2007\)](#); [Von Gaudecker \(2015\)](#)) or take priced risk (e.g., [Ibert et al. \(2017\)](#); [Ben Naim and Sokolinski \(2017\)](#)). 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. Such households can either invest more than their optimal share or simply not participate in wealth management. In the latter case, reducing the minimum relaxes funding 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.

We find strong support for the previous intuition, in that the model quantitatively replicates our key empirical findings. Thus, from a positive perspective, the large estimated effect of the reduction stems from households' rational response to differences in portfolio characteristics. Notably, we explain this effect without taking a stance on the magnitude of difficult-to-quantify psychological costs (e.g., [Gennaioli et al. \(2015\)](#)).

Turning to welfare, the reduction raises middle-class households' welfare by 2%, based on the standard lifetime consumption metric, while it has almost zero effect on the upper class. For reference, this welfare gain equates to that of a permanent 4 pps increase in the equity premium with no reduction. 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 22% (0.4 pps) more than those under age 35. Intuitively, 83% of over-55 households would never have participated in wealth management without the reduction, being permanently constrained. By contrast, 77% of under-35 households would have eventually overcome the previous minimum, being temporarily constrained. Therefore, contrary to their popular image as the investment of choice for millennials, we find that robo advisors add significant value to households at a later stage in life.

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 households. This policy conclusion comes with the caveat that we study a period of rapid growth in the market for

robo advising (e.g., [D’Acunto and Rossi 2020](#)), and so the same research design may yield different estimates in other economies.

We conclude our introduction 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 literatures. First, we contribute to a literature on how new technologies affect the financial sector. Within this FinTech literature, we contribute most specifically to a nascent literature on robo advisors by showing how they democratize wealth management. Prior and contemporaneous work examines how robo advisors affect wealthier households. We differ not only in our focus on the less-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 Brière \(2020\)](#); [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, we are the first to structurally evaluate the effects of robo advisors. Like the existing research, however, our findings contrast with the high fees, underperformance, and misaligned incentives associated with traditional asset managers and financial advisors (e.g., [French 2008](#); [Bailey, Kumar and Ng 2011](#); [Christoffersen, Evans and Musto 2013](#); [Del Guercio and Reuter 2014](#); [Linnainmaa, Melzer and Previtro 2021](#); [Linnainmaa et al. 2020](#)).

More broadly within the FinTech literature, we show how FinTech affects financial inclusion and wealth inequality in a new setting. In terms of financial inclusion, 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

⁴It is well-known that financial returns increase by wealth and education, and our findings suggest that FinTech can

Second, we contribute theoretically and empirically to the household finance literature. 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 self-managed versus professionally managed portfolios, and, like [Fagereng et al. \(2017\)](#), calibrate it using portfolio-level data. These features allow us to parsimoniously match the data without 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 minimum investments 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 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 minimum investments 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 services beyond simply advice or portfolio management.

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

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 perspective differs from that of models predicting a positive relationship between FinTech and inequality (e.g., [Begenau, Farboodi and Veldkamp 2018](#); [Kacperczyk, Nosal and Stevens 2019](#); [Mihet 2020](#)).

per-portfolio management costs relative to the traditional approach of manually constructing and managing a client’s portfolio. 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. Note that Wealthfront stands out as the only robo advisor that relies purely on automation, with no option for a human advisor.

2.2 Robo Portfolios

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 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

⁵Strictly 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).

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, which represents a sizeable decline from the standpoint of most U.S. households. For reference, \$5,000 equals 30% of the median household's liquid assets (\$17,000), and it defines the 37th percentile of the U.S. wealth distribution, according to the 2016 Survey of Consumer Finances. Prior to the reduction, therefore, half of U.S. households could not participate with the advisor 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 zero for accounts under \$10,000 or 0.25 pps for larger accounts, the reduction was not intended to increase 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, Betterment, 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

⁶In 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."

⁷This advisor 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).

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 because of the reduction affects such inequality, but 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 10 contains details. For the rest of the paper, we use the term “robo participant” to describe households who have invested money with the robo advisor, Wealthfront.

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 reduction in minimum, and it 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, and so 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; householder 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 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, and, thus, we can credibly restrict the set of non-robo portfolios to those managed by the portfolio’s owner (i.e., self-managed). This feature enables us to evaluate the choice between managing one’s own portfolio versus delegating it to a professional under realistic parameter values. Second, we observe pairs of self-managed and counterfactual robo portfolios for both robo participants and non-participants, with the latter comprising 55% of the sample. This feature also enables a realistic parameterization, as it avoids selection bias from, say, the possibility that only households who cannot invest efficiently on their own delegate to a professional. By contrast, selection from the fact that the dataset only includes households who consult the online tool does not pose a concern, since these households are on the margin of robo participation and, thus, exactly those whose behavior we seek to model.

Comparing the roles of the two datasets: the deposits dataset enables our main empirical analysis, and the portfolio dataset enables the model that explains our empirical analysis. Thus, each dataset accomplishes a separate goal, and so we do not need to combine them. Furthermore, such a combination would necessitate a fuzzy merge because, out of respect for client privacy, we only observe a household’s demographic characteristics but not a unique, cross-dataset identifier.

⁸For 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).

Therefore, we are careful to remark that our welfare statements in Section 8 only pertain to a rational household who, per the quality of the model’s fit in Section 7, behaves in the same way as the households in our deposits dataset.

3.3 Auxiliary Datasets

Our analysis relies on a variety of auxiliary datasets, which we discuss in turn. Of these datasets, the most important 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. We defer summarizing the portfolio dataset until we describe the model in Section 6. The upper panel of Table 1 compares households who become robo participants after the reduction in account minimum with existing participants. The upper panel shows how households who become robo participants after the reduction (i.e., 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.

The lower panel 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 obser-

vation in Sections 4.3 and 8.2.

Lastly, new middle-class robo participants exhibit persistent investment behavior, both in absolute terms and relative to existing middle-class and new upper-class participants. For example, 97% do not close their account over our sample period, mirroring the 98% non-closure rate among new upper-class participants. Additionally, 72% of new middle-class robo participants make a subsequent deposit, comparable to 71% among new upper-class participants. Such subsequent deposit-making resembles the “dollar cost averaging” strategy commonly advocated by practitioners, which Brennan et al. (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 relaxing funding 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 ex-

actly 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 investment 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 behavior between the middle versus upper classes represents the effect of funding 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} ,

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

where i indexes household; \mathcal{T} indexes the pre-reduction period (i.e., $\mathcal{T} = 0$) vs. 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 higher 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.⁹

The additional terms in equation (1) capture channels distinct from the funding 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 investment constraints because the minimum changes over time. The interaction between X_i and $Post_T$ 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 investment constraints, then the coefficient ψ would separately capture this effect.

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

$$\Delta Participant_i \equiv New Participant_i = \mu Middle_i + \psi X_i + \varrho + u_i, \quad (2)$$

where $New Participant_i$ indicates if household i becomes a robo participant after the reduction; and $u_i \equiv \Delta v_i$. We estimate equation (2) on the set of eventual robo participants, and, therefore, μ equals the reduction’s effect on the probability of robo participation, conditional on eventually participating. Section 4.4 shows explicitly how this statistic quantifies the democratization of the robo market.

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

$$0 = \mathbb{E} [Middle_i \times u_i | X_i]. \quad (3)$$

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,

⁹Explicitly, 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.

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 Vanguard and Charles Schwab. 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 funding constraint (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, per Appendix Table A2. We also find that the effects of the reduction are stronger for risk-averse households, 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 11, 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 that exceed \$8,400 ($2 \times 0.1 \times 42,000$). Columns (6) and (7) show that the estimates based on these alternative measures of $Middle_i$ equals 0.10 and 0.16. This range straddles our baseline estimate of 0.14, suggesting that it is not biased because of measurement error.

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 investment constraints and, thus, $\mu = 0$. Appendix 11 shows how the growth rate under this counterfactual equals

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

Our statistic of interest is

$$\eta \equiv g - g^C, \quad (5)$$

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 11. Interpreting the first row, the baseline estimates from Table 2 imply that the reduction increases the overall number of robo participants by 14%, which is driven by a 110% increase in the number of middle-class participants. In relation to Table 2, the 110% 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. The additional estimates in Table 3 imply an increase in the number of middle-class participants between 56% and 129%.

The large effects in Table 3 suggest that households have significant demand for professional

wealth management, but many face funding constraints. In Section 6, we propose a theory that can quantitatively explain this demand.

5 Robustness

We assess the internal validity of the baseline results. Specifically, we directly assess the funding 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

5.1.1 Regression Evidence on Bunching

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.

Table 4 reports the results. Column (1) shows that new middle-class participants are 30 pps more likely to invest under \$5,000 than new upper-class participants, which matches the graphical evidence from panel (a) of Figure 4. Reiterating the discussion from Section 4.1, 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 funding constraints are plausibly more binding (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. This effect is also 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 constraints.

5.1.2 Effect on Robo Share

If new middle-class participants increase the share of their portfolio allocated to the robo advisor more than new upper-class participants, then they plausibly experience a relaxation of funding constraints. To see why, consider a null hypothesis in which the reduction has no effect on funding constraints. Under this null hypothesis, both middle and upper-class participants could have already invested their unconstrained-optimal robo share before the reduction. Consequently, any increase in robo share reflects an investment for non-financial purposes (e.g., excitement), and, for the sake of argument, non-financial motivations do not vary by wealth.¹⁰ Thus, the null hypothesis predicts a uniform increase in robo share by wealth class.

We test this null hypothesis by estimating equation (2) on the set of new robo participants, after replacing the outcome with a participant's change in robo share. Define the change in robo share as

$$\Delta Robo Share_i = \frac{Robo Investment_i}{Liquid Assets_i}, \quad (6)$$

where $Robo Investment_i$ equals the value of net deposits by i in the post-reduction period. The results in column (5) of Table 4 show that new middle-class robo participants increase their risky share by 13.5 pps relative to their upper-class peers. In column (6), we again find evidence of substantial heterogeneity within the middle class. In particular, new middle-class households from the second quintile increase their robo share 27 pps more than the upper class, compared to 12 pps for those from the third quintile. Together, these results lead us to reject the null hypothesis that the middle class experiences no relaxation of funding 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

$$New Participant_{i,t} = \mu (Middle_i \times Post_t) + \zeta_i + \varrho_t + u_{i,t}, \quad (7)$$

¹⁰Alternatively, any increase in robo share reflects an investment made for financial purposes after learning about the advisor through advertising. Section 5.2 shows how advertising affects the middle and upper classes equally, and so this alternative possibility also predicts a uniform increase in robo share by wealth class.

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. As a first step, we estimate equation (7) as-is and report the results in column (1) of Table 5. 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 (32×0.007), which is on par with the estimated effect in Table 2.

5.2.1 Media Attention and Targeted Advertising

We examine whether media attention, advertising, or other changes in visibility affect our results 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: *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 *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.

Our primary concern is that media attention and advertising around the reduction disproportionately influence middle-class households. We address this concern by interacting *Middle_i* with the previous two proxies, thus allowing middle-class households to respond differentially to changes in the robo advisor's visibility. The corresponding coefficient of interest in columns (2) and (3) 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

More generally, our baseline results may confound any dynamic effect that occurs over our observation window and disproportionately affects the middle class. Examples of such effects include a secular trend in middle-class households' demand for automated asset management or changes in industry competition for the middle class.

We address this concern by replacing $Post_t$ in equation (7) with a set of indicator variables that equal one if the month of week t is k months before the reduction, denoted *Months Before_{t,k}*, or after it, denoted *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., $Months\ Before_{t,1}$).

The results in column (4) 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 $Months\ Before_{t,k}$. 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.

5.2.3 Dynamic Response by Millennials

Commentators often portray automated asset management as the investment of choice for millennials, defined as being born between 1981 and 1996. If millennials respond more strongly to an increase in the robo advisor's visibility because of, say, technological savviness, then our baseline results suffer upward bias because millennials also disproportionately belong to the middle class. We assess the scope for such bias by interacting $Post_t$ with an indicator for whether household i is a millennial. The highly insignificant coefficient on this interaction in column (5) of Table 5 implies that millennials do not respond to the reduction for reasons apart from belonging to the middle class.

5.3 Business Stealing and Gambling

As described in Appendix 11, 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 on household investment, as 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, solution, and calibration 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 relative quantities of compensated and uncompensated risk in these two portfolios. We instead let the data inform these quantities, as described in Section 6.2. Second, the two portfolios differ in that the latter requires a minimum investment, M . We soon narrow our focus to the particular, automated wealth manager described in Section 3. 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 minimum investment, that is, the robo portfolio.

6.1.1 Preferences

As in our empirical analysis, let i index household. Time is discrete, and household i reoptimizes in each year t . 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 \bar{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

$$U_{i,t} = \mathbb{E}_t \left[\sum_{\tau=t}^{\tau=\bar{T}} \delta^{\tau-t} \left(\prod_{j=t}^{j=\tau-1} p_j \right) \frac{C_{i,\tau}^{1-\gamma}}{1-\gamma} \right], \quad (8)$$

where δ is the discount factor. The absence of a bequest motive in equation (8) 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 \(1991\)](#)). These resources consist of 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 $R_{i,t}^S$ and can be likened to a discount brokerage account; and a risky portfolio overseen by an automated wealth manager, which gives return $R_{i,t}^A$. The last of these portfolios requires a minimum investment of M and a 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 the empirical covariance between the two portfolios discourages such simultaneous holding.

Following convention in the household finance literature, we introduce a factor structure for risky returns (e.g., [Calvet, Campbell and Sodini \(2007\)](#); [Von Gaudecker \(2015\)](#)). This approach addresses the well-known challenge of estimating expected returns in finite samples (e.g., [Merton \(1980\)](#)), which would be particularly problematic in our setting given the limited history of the assets in our portfolio-level dataset. Explicitly, we suppose the return on portfolio $\mathcal{P} \in \{S, A\}$ evolves according to

$$R_{i,t}^{\mathcal{P}} = \beta_i^{\mathcal{P}} F_t + \epsilon_{i,t}^{\mathcal{P}}, \quad (9)$$

where F_t is a vector of priced risk factors, normally distributed with mean π^F and covariance matrix Σ^F ; $\beta_i^{\mathcal{P}}$ is the loading of portfolio \mathcal{P} on this vector for household i ; and $\epsilon_{i,t}^{\mathcal{P}}$ is an idiosyncratic shock, normally distributed with mean zero and volatility of $\sigma_{\epsilon,i}^{\mathcal{P}}$. The quantity of compensated risk, $\beta_i^{\mathcal{P}}$, 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, $\sigma_{\epsilon,i}^{\mathcal{P}}$, may vary across both portfolios and investors.

As summarized by [Gomes \(2020\)](#), workhorse models typically feature a fixed dollar cost of stock market participation. Without such a cost, these models generally predict that almost all households participate in risky asset markets, whereas Appendix Table A3 shows how, in reality,

only 37% do so. We can dispense with this additional parameter by incorporating both imperfect diversification, per equation (9), and a minimum investment, M . Intuitively, households seek professional management because they struggle to diversify on their own, as documented shortly. However, accessing professional management requires such a large minimum investment that many households decide not to participate in the stock market at all. Thus, while our model aims to explain participation in wealth management, as distinct from stock market participation, it still gives a fairly accurate prediction that only 44% of all households participate in the stock market.

6.1.3 Labor Income

Households retire at age \underline{T} . 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

$$\log(Y_{i,t}) = f_t + \tilde{\zeta}_{i,t} + v_{i,t}, \quad (10)$$

where f_t is a deterministic function of age, t , as distinct from the year; $\tilde{\zeta}_{i,t}$ is a transitory shock, normally distributed with mean zero and volatility of $\sigma_{\tilde{\zeta}}$; and $v_{i,t}$ is a permanent shock that evolves according to

$$v_{i,t} = v_{i,t-1} + \underbrace{\Xi_t + \omega_{i,t}}_{v_{i,t}}, \quad (11)$$

where $v_{i,t}$ is normally distributed with mean zero and volatility of σ_v . Equation (11) 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 papers 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 (10).

Our empirical analysis primarily concerns investment prior to retirement, and so we model

the post-retirement period more simply than do workhorse models. Households do not receive labor income for $t > \underline{T}$. Thus, for $t = \underline{T} + 1, \dots, \bar{T}$, 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 12 states the problem formally. A household's expected lifetime utility as of age $\underline{T} + 1$ then has the familiar form

$$V_{\underline{T}}(W_{i,\underline{T}+1}) = B \frac{W_{i,\underline{T}+1}^{1-\gamma}}{1-\gamma}, \quad (12)$$

where B is a function of parameters.¹¹ One can interpret equation (12) as a microfounded bequest motive.

6.1.4 Consolidated Problem

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

$$C_{i,t} = [1 - \alpha_{i,t}^f - \alpha_{i,t}^S - \alpha_{i,t}^A] W_{i,t}. \quad (13)$$

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

$$\alpha_{i,t}^f \geq 0, \quad \alpha_{i,t}^S \geq 0, \quad \alpha_{i,t}^A \geq 0, \quad (14)$$

$$1 - \alpha_{i,t}^f - \alpha_{i,t}^S - \alpha_{i,t}^A \geq 0, \quad (15)$$

$$\alpha_{i,t}^A = 0 \quad \text{or} \quad \alpha_{i,t}^A \geq \frac{M}{W_{i,t}}. \quad (16)$$

Constraint (14) rules out borrowing and shorting, which we subsequently relax in Section 8.1.3. Constraint (15) ensures nonnegative consumption. Both of these constraints are standard. The third constraint, (16), requires a minimum investment of M to participate in wealth management.

¹²

As noted by [Cocco, Gomes and Maenhout \(2005\)](#), this setup leads to a problem with two state

¹¹Explicitly, $B = \sum_{\tau=\underline{T}+1}^{\bar{T}} \delta^{\tau-\underline{T}-1} \left[\frac{1-\chi}{1-\chi^{\tau-\underline{T}}} \right] \left[\delta(1+R^f) \right]^{\frac{\tau-\underline{T}-1}{\gamma}}$ with $\chi = \delta^{\frac{1}{\gamma}} (1+R^f)^{\frac{1-\gamma}{\gamma}}$.

¹²Technically, constraint (16) reflects a minimum balance constraint. This simplification does not affect the results, and it substantially reduces the model's computational intensity because it eliminates an additional state variable. See Appendix 12.

variables, age (t) and liquid assets ($W_{i,t}$). The latter evolves according to

$$W_{i,t+1} = \left[\alpha_{i,t}^f (1 + R^f) + \alpha_{i,t}^S (1 + R_{i,t+1}^S) + \alpha_{i,t}^A (1 + R_{i,t+1}^A) \right] W_{i,t} + Y_{i,t+1}. \quad (17)$$

Intuitively, the problem is homogeneous in permanent labor income, $v_{i,t}$, allowing us to effectively remove it from the state space. Collectively, therefore, household i of age t solves the following Bellman equation,

$$V_t(W_{i,t}) = \max_{\alpha_{i,t}^f, \alpha_{i,t}^S, \alpha_{i,t}^A} \left\{ \frac{C_{i,t}^{1-\gamma}}{1-\gamma} + \delta p_t \mathbb{E}_t [V_{t+1}(W_{i,t+1})] \right\} \quad (18)$$

s.t. (9)-(17).

We solve equation (18) using standard numerical methods described in Appendix 12. Briefly, we: discretize the state space defined by age and liquid assets; solve equation (18) for age \underline{T} ; and iteratively solve equation (18) 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 12 contains details.

6.2.1 Portfolio Parameters

We use the portfolio dataset described in Section 3.2 to realistically calibrate the vector of portfolio parameters, $\{\sigma_{\epsilon,i}^S, \sigma_{\epsilon,i}^A, \beta_i^S, \beta_i^A\}$. We reiterate that this dataset includes security-level information on self-managed portfolios for households on the margin of participating with the robo advisor. This sampling restriction is ideal, since these are exactly the households whose behavior we seek to explain.

Our calibration proceeds in three steps. First, we estimate the parameter vector $\{\sigma_{\epsilon,j}^S, \sigma_{\epsilon,j}^A, \beta_j^S, \beta_j^A\}$ for each of the 1,913 pairs of self-managed and robo portfolios in the dataset, which we index by j . For a given vector of risk factors, F , we estimate the following pricing equation for each security k in the portfolio dataset,

$$R_{k,m} = \beta_k F_m + \epsilon_{k,m}, \quad (19)$$

where m indexes month; and $R_{k,m}$ denotes the monthly return on security k in excess of the risk-free return. Let w_j^S denote a vector of weights across securities k for the self-managed portfolio j , and, likewise, let w_j^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

$$\sigma_{\epsilon,j}^{\mathcal{P}} = \sqrt{w_j^{\mathcal{P}'} \hat{\Sigma}_\epsilon w_j^{\mathcal{P}}}, \quad \beta_j^{\mathcal{P}} = w_j^{\mathcal{P}'} \hat{\beta}, \quad (20)$$

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

Second, we specify the pricing vector, 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 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 A4](#) shows how robo portfolios feature 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 A5](#). For reference, the mean and volatility of the market factor equal 7.6% and 14.7%, respectively.

Third, 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 a fifth-order polynomial in the portfolio holder's age and the holder's log liquid assets. Then, we substitute the fitted values into [equation \(9\)](#) according to the age and liquid assets of household i . We find similar results when simply substituting the sample average into [equation \(9\)](#), largely because self-managed portfolios vary little according to the holder's demographic characteristics.

[Table 7](#) provides intuition for our coming results by summarizing the portfolio dataset. Robo portfolios feature a 30 pps higher Sharpe ratio, which, interestingly, does not differ between the middle and upper classes. This higher Sharpe ratio partly reflects a 2 pps higher expected return on robo portfolios. By construction, a higher expected return reflects greater exposure to priced risk. However, robo portfolios contain less total risk because they are better diversified, with 11 pps lower idiosyncratic volatility. [Section 8](#) quantifies the relative welfare impact of priced risk and diversification.

We reiterate that 55% of households in the dataset never become robo participants. If, by con-

trast, the dataset only included eventual participants, then Table 7 might overstate the differences between self-managed and robo portfolios because, say, eventual participants struggle more with self-management. Therefore, the portfolio dataset is ideal for the purpose of calibration.

6.2.2 Other Parameters

We choose preference parameter values of $\gamma = 9$ and $\delta = 0.96$, consistent with the literature (e.g., Cocco et al. (2005); Fagereng et al. (2017)). Appendix Table A8 shows how we obtain similar values when structurally estimating these parameters, as described in Appendix 12. We follow Cocco et al. (2005) in our calibration of labor income parameters. Accordingly, the deterministic component of income, f_t , is a third-order polynomial in householder age, and the coefficients equal those estimated by Cocco et al. (2005) for their baseline analysis. Similarly, we parameterize $\sigma_v = 0.103$, $\sigma_\xi = 0.271$, and $\beta^Y = 0.001$, all of which follow from the estimates in Cocco et al. (2005). We parameterize $\phi = 0.001$, corresponding to 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 $\bar{T} = 100$, 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

The model has positive implications insofar as it can explain the quasi-experimental evidence in Section 4. We assess these implications by calculating 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.¹³ By construction, the theoretical effect works through a relaxation of constraints, and, as we confirmed in Section 5, so does the empirical effect.

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

¹³The 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, unlike related papers that study how a representative household responds to recurring shocks over her life.

liquid assets, both for participants who optimally invest under the previous minimum (Existing Participants) and for those who optimally invest under the reduced minimum (New Participants). 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 13 shows how we obtain similar quality of fit under different pricing factors.

Table 8 reproduces three other sets of statistics. First, panel (a) replicates Figure 5 in terms of wealth classes, as opposed to wealth quintiles. As with Figure 5, the theoretical distribution matches the empirical distribution quite closely prior to the reduction. The model predicts a large 8.6 pps increase in the share of middle-class participants. While less than the 15 pps increase observed in the data, this error does not pose a concern given the model’s parsimony and ability to match other statistics. For example, the model predicts the robo share of new middle and upper-class participants remarkably well, as shown in panel (b). This finding is unexpected because we do not target robo share in our calibration. Lastly, panel (c) reproduces the growth in the number of robo participants reported in Table 3. The model matches the overall growth rate well. It understates the growth in middle-class participation, but, for the same reason as in panel (a), this does not pose a concern.

The quality of fit in Figure 5 and Table 8 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 the self-managed and robo portfolios had the same return, $R_{i,t}^A = R_{i,t}^S$, then the model would counterfactually predict no change in participation. Therefore, middle-class households’ demand for wealth management must fundamentally stem from differences in diversification (i.e., $\sigma_{\epsilon,i}^P$) or priced risk (i.e., β_i^P). We decompose the contribution of these two channels in the next section. For now, we conclude that households act according to a rational model with frictions that keep them from forming portfolios as a robo advisor would, without the need for additional parameters that govern “trust” (e.g., Guiso et al. (2008)) or “peace-of-mind” (e.g., Gennaioli et al. (2015)).

8 Welfare Implications

According to the logic of revealed preference, the reduction improves new robo participants’ welfare. We focus on the channels through which the reduction improves welfare and the distribution of this gain across households. As standard, we measure household i ’s welfare gain by

the percent increase in annual consumption under the previous minimum that raises her expected lifetime utility by the same amount as the reduction. Letting \underline{V}_i and \bar{V}_i denote the value functions under the minimums of \$5,000 and \$500, respectively, this welfare gain equals

$$q_i = \left(\frac{\bar{V}_i}{\underline{V}_i} \right)^{\frac{1}{1-\gamma}} - 1, \quad (21)$$

as shown in Appendix 12.

Like other papers that use life cycle models to study distributional effects (e.g., [Gete and Zecchetto \(2018\)](#)), we recalculate equation (21) separately for subpopulations defined by wealth and age. Then, we decompose equation (21) into three terms that reflect the particular gain from changes in diversification, priced risk, and risky share. Explicitly,

$$q_i = \underbrace{\left[\left(\frac{\bar{V}_i |_{\alpha, \sigma_{\epsilon,i}^P}}{\underline{V}_i} \right)^{\frac{1}{1-\gamma}} - 1 \right]}_{\text{PricedRisk}} + \underbrace{\left[\left(\frac{\bar{V}_i |_{\alpha}}{\underline{V}_i} \right)^{\frac{1}{1-\gamma}} - \left(\frac{\bar{V}_i |_{\alpha, \sigma_{\epsilon,i}^P}}{\underline{V}_i} \right)^{\frac{1}{1-\gamma}} \right]}_{\text{Diversification}} + \underbrace{\left[\left(\frac{\bar{V}_i}{\underline{V}_i} \right)^{\frac{1}{1-\gamma}} - \left(\frac{\bar{V}_i |_{\alpha}}{\underline{V}_i} \right)^{\frac{1}{1-\gamma}} \right]}_{\text{RiskyShare}} \quad (22)$$

where $\alpha_i \equiv \alpha_i^S + \alpha_i^A$ is household i 's risky share; $\bar{V}_i |_{\alpha}$ 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, $\bar{V}_i |_{\alpha, \sigma_{\epsilon,i}^P}$ 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 (22) 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., $\beta_i^A \neq \beta_i^S$). The second term equals the gain when the two portfolios also differ in idiosyncratic risk (i.e., $\sigma_{\epsilon,i}^A \neq \sigma_{\epsilon,i}^S$), such that portfolio returns equal their actual values according to equation (9), but households still cannot increase their risky share. Notably, these first two terms equals zero for households who do not already 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 middle-class participant gains 2% in lifetime consumption, compared to almost nothing for the average upper-class participant.

¹⁴For reference, the workhorse models referenced earlier generally consider a one percent gain in lifetime consumption economically significant. To place a 2% gain in perspective, panel (d) shows how lifetime consumption rises by 1.7% for the middle-class participants in panel (a) following a permanent 4 pps increase in the equity premium, holding the minimum fixed at \$5,000. Likewise, a one standard deviation increase in log labor income raises their lifetime consumption by 1.4%.

8.1.1 Decomposition of Channels

Panel (b) of Table 9 decomposes the total welfare gain according to equation (22). 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. First, the gain from exposure to priced risk matches empirical evidence that fund managers are compensated for taking such risk, cited in the introduction. 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 predominantly improves welfare through diversification.

Lastly, the gain from a higher risky share principally reflects the gain from becoming a stock market participant. Workhorse models predict that almost all households participate in the stock market. In our model, however, only 44% participate in the stock market, which is close to the actual share of 37% shown in Appendix Table A3. Intuitively, households seek professional management because they struggle to diversify on their own. 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, given the choice between an underdiversified, self-managed portfolio with a reasonable risky share or a well-diversified, professionally managed portfolio with an excessive risky share, 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.

¹⁴The 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 minimum requires an investment that exceeds their unconstrained-optimum.

8.1.2 Model Extensions

Panel (c) summarizes welfare gains under three model extensions, the details of which are in Appendix 12. First, middle-class households experience a 2.5% welfare gain when we introduce a per-period cost of participating in the self-managed portfolio. The per-period cost equals \$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 (Gomes (2020)). We account for explicitly for these effects, and so we choose a smaller value that, say, captures time costs associated with rebalancing. That welfare gains increase under this model extension suggests that our baseline model produces conservative results.

Next, the middle-class welfare gain equals 1.7% when 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 minimum investment. Consequently, reducing the minimum has a smaller relative effect on lifetime utility and, thus, we obtain a 0.3 pps smaller welfare gain than the baseline gain of 2%.

Lastly, we find a 3.3% middle-class welfare gain when incorporating 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 value of assets that households can liquidate at any time. Together, these effects make a lower minimum investment more valuable, leading to a 1.3 pps larger welfare gain than in the baseline model.

8.1.3 Robustness

Appendix Table A7 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 10 repeats the core exercise in Table 9 by age. Interestingly, the average welfare gain for middle-class participants increases in age, as shown in panel (a). This finding is surprising given that robo advisors claim to "build our products and services for millennials" (e.g., Hutchins

(2020)). Intuitively, many young non-participants would have eventually become robo participants even without the reduction.

In more detail, we first note that middle-class robo participation increases in age, as shown explicitly in Appendix Figure A4. This finding corroborates a well-known result within the life cycle literature. Paraphrasing Fagereng et al. (2017), it reflects how, within the middle class, 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% probability of eventually participating by retirement, as shown in panel (b).¹⁵ By contrast, households over age 55 who have not yet participated with the robo have only a 17% probability of eventually participating.

The reduction uniformly increases middle-class robo participation across age, as we found empirically in Table 2 and show theoretically in Appendix Figure A4. Consequently, it raises the eventual probability of participation by 74 pps for households over age 55 but by only 23 pps for those under age 36. Thus, the reduction relaxes a “temporary” constraint on the younger middle class but a “permanent” constraint on the older middle class, and so the latter gains more,

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, without the need for government intervention. This conclusion comes from studying a large and unexpected reduction in minimum investment 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 set of investment opportunities available to the modestly wealthy (e.g., myRA, OregonSaves, NEST).

Second, from the perspective of economic theory, our results support models in which households behave rationally given limits on their ability to invest efficiently. We arrive at this conclu-

¹⁵The probability equals $P_t^R = \sum_{\tau=t+1}^{\tau=T} \Delta P_t \prod_{j=t+1}^{j=\tau-1} (1 - \Delta P_j)$, 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, $P^R(T) = 0$ because a non-participant of age T retires in the following year.

sion 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. 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 hope that future work can study whether these results extend to aging or developing economies.

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
- Bailey, W., Kumar, A. and Ng, D.: 2011, Behavioral Biases of Mutual Fund Investors, *Journal of Financial Economics* . 1
- Barber, B. M., Huang, X., Odean, T. and Schwarz, C.: 2020, Attention Induced Trading and Returns: Evidence from Robinhood Users. 1
- Barberis, N., Huang, M. and Thaler, R. H.: 2006, Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing, *American Economic Review* . 1
- Bartlett, R., Morse, A., Stanton, R. and Wallace, N.: 2021, Consumer-Lending Discrimination in the FinTech Era, *Journal of Financial Economics* . 1
- Begenau, J., Farboodi, M. and Veldkamp, L.: 2018, Big Data in Finance and the Growth of Large Firms, *Journal of Monetary Economics* . 4
- Ben-David, I., Franzoni, F., Kim, B. and Moussawi, R.: 2021, Competition for Attention in the ETF Space, *NBER Working Paper* . 1
- Ben Naim, G. and Sokolinski, S.: 2017, What Does Compensation of Portfolio Managers Tell Us about Mutual Fund Industry? Evidence from Israeli Tax Records. 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 Bri re, 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, 11.4
- 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, 10.3.1
- 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, 6.1.2, 13
- 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
- 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
- Christoffersen, S., Evans, R. and Musto, D.: 2013, What Do Consumers Fund Flows Maximize? Evidence from Their Brokers Incentives, *The Journal of Finance* . 1
- 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, 6.1.4, 6.2.2, 6, 12.1.2
- 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* . 12.1.1
- 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

- Del Guercio, D. and Reuter, J.: 2014, Mutual Fund Performance and the Incentive to Generate Alpha, *Journal of Finance* . 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, 6.2.2, 8.2
- Fagereng, A., Guiso, L., Malacrinò, 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
- French, K. R.: 2008, Presidential Address: The Cost of Active Investing, *The Journal of Finance* . 1
- Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T. and Walther, A.: 2021, Predictably Unequal? The Effects of Machine Learning on Credit Markets, *The Journal of Finance* . 1
- Fuster, A., Plosser, M., Schnabl, P. and Vickery, J.: 2019, The Role of Technology in Mortgage Lending, *Review of Financial Studies* . 1
- Gennaioli, N., Shleifer, A. and Vishny, R.: 2015, Money Doctors, *Journal of Finance* . 1, 7
- Gete, P. and Zecchetto, F.: 2018, Distributional Implications of Government Guarantees in Mortgage Markets, *Review of Financial Studies* . 8, 12.2
- Gomes, F.: 2020, Portfolio Choice Over the Life Cycle: A Survey, *Annual Review of Financial Economics* . 1, 6.1.2, 8.1.2
- 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, 7
- Guvenen, F., Ozkan, S. and Song, J.: 2014, The Nature of Countercyclical Income Risk. 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
- 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

- Ibert, M., Kaniel, R., Van Nieuwerburgh, S. and Vestman, R.: 2017, Are Mutual Fund Managers Paid for Investment Skill?, *Review of Financial Studies* . 1
- Kacperczyk, M., Nosal, J. and Stevens, L.: 2019, Investor Sophistication and Capital Income Inequality, *Journal of Monetary Economics* . 4
- Kalda, A., Loos, B., Previtero, A. and Hackethal, A.: 2021, Smart(Phone) Investing? A within Investor-time Analysis of New Technologies and Trading Behavior., *NBER Working Paper* . 1
- Kaniel, R. and Parham, R.: 2017, WSJ Category Kings â The Impact of Media Attention on Consumer and Mutual Fund Investment Decisions, *Journal of Financial Economics* . 1, 4.2
- Linnainmaa, J. T., Melzer, B. T. and Previtero, A.: 2021, The Misguided Beliefs of Financial Advisors, *The Journal of Finance* . 1
- Linnainmaa, J. T., Melzer, B. T., Previtero, A. and Foerster, S.: 2020, Investor Protections and Stock Market Participation: An Evaluation of Financial Advisor Oversight. 1
- 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* . 2.3
- 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* .
- Tauchen, G. and Hussey, R.: 1991, Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models, *Econometrica* . 12.1.2


- 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, 8.1.2
- Von Gaudecker, H.-M.: 2015, How Does Household Portfolio Diversification Vary with Financial Literacy and Financial Advice?, *The Journal of Finance* . 1, 6.1.2
- Welch, I.: 2020, The Wisdom of the Robinhood Crowd, *NBER Working Paper* . 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, 10.3.3

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 minimum investment (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



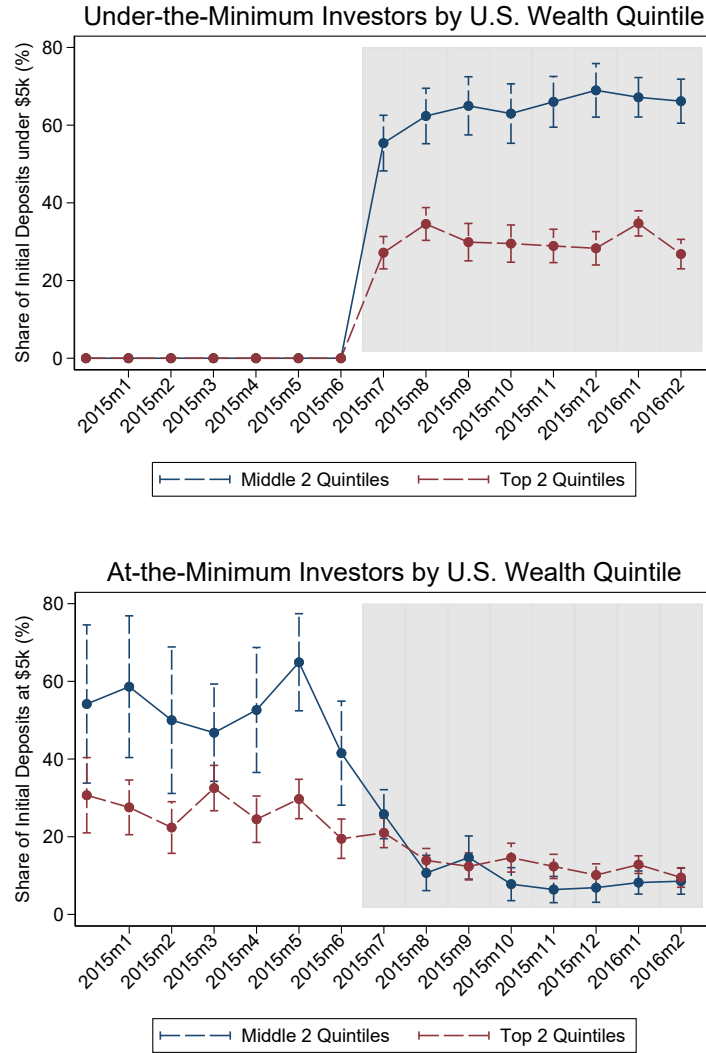
Note: This figure plots the share of robo participants from each quintile of the U.S. wealth distribution. The share is calculated separately for households who participated before the reduction in minimum investment (Existing Participants) and who become participants after the reduction (New Participants). 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 Group




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. 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 minimum investment. Brackets are 95% confidence intervals.

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 minimum investment. Brackets are 95% confidence intervals.

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 among participants who find it optimal to participate under the previous minimum (Existing Participants) and among 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

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

Number of Existing Participants: 4,366

Number of New Participants: 5,336

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). 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 lower panel summarizes households from the second or third U.S. wealth quintile. Appendix 10 contains additional variable descriptions.

Table 2: Democratization of the Robo Market after the Reduction

y =	<i>New Participant_i</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Middle_i</i>	0.219 (0.000)	0.151 (0.000)	0.137 (0.000)	0.180 (0.000)	0.126 (0.000)	0.095 (0.001)	0.155 (0.000)
<i>Middle_i × Age_i</i>				-0.001 (0.384)			
<i>Middle_i × Risk Averse_i</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	0.088	0.098
Number of Observations	9,349	9,349	9,349	9,349	9,349	9,349	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

$$New\ Participant_i = \beta Middle_i + \delta X_i + \tau + u_i,$$

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 (4)-(6) assess the scope for measurement error by remeasuring $Middle_i$ using alternative measures: an indicator for whether i belongs to the second U.S. wealth quintile (Second Wealth Quintile), which we later denote in Table ?? as *Lower Middle_i*; 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 U.S. wealth quintile (Middle with Buffer); and an indicator for whether i belongs to the second or third U.S. wealth quintile, after assigning a missing value to households whose reported wealth quintile does not equal their imputed wealth quintile based on the boosted trees imputation methodology described in Section 5.1 applied to the SCF dataset (Predictable Middle). The sample consists all robo participants in the Wealthfront 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 log ($Income_i$), Age_i , and $Risk\ Averse_i$, 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)
<u>Baseline Estimates:</u>						
Table 2, Column (3)	119.4%	106.0%	13.4%	239.4%	131.6%	107.8%
<u>Additional Estimates:</u>						
Table 2, Column (6)	119.4%	118.4%	1%	166.7%	112.7%	53.9%
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. 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: Effects on Funding Constraints and Robo Share

$y =$	<i>Under Minimum_i</i>		<i>At Minimum_i</i>		<i>Robo Share_i</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Middle_i</i>	0.294 (0.000)		0.253 (0.000)		0.135 (0.000)	
<i>Second Quintile_i</i>		0.555 (0.000)		0.309 (0.006)		0.272 (0.000)
<i>Third Quintile_i</i>		0.269 (0.000)		0.248 (0.000)		0.122 (0.000)
<i>Middle_i × New Participant_i</i>			-0.316 (0.000)			
<i>Second Quintile_i × New Participant_i</i>				-0.467 (0.000)		
<i>Third Quintile_i × New Participant_i</i>				-0.302 (0.000)		
<i>New Participant_i</i>			-0.149 (0.000)	-0.149 (0.000)		
Controls	Yes	Yes	No	Yes	Yes	Yes
State FE	Yes	Yes	No	No	Yes	Yes
R-squared	0.156	0.165	0.096	0.097	0.087	0.098
Number of Observations	5,088	5,088	6,890	6,890	5,088	5,088

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

$$Y_i = \lambda_0 \text{Middle}_i + \lambda_1 X_i + \lambda_2 + v_i,$$

where Y_i is a measure of constraints imposed on i by the previous account minimum. Columns (1)-(2) 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. Column (2) tests for a change in bunching behavior by middle-class participants by including the interaction between *Middle_i* and *New Participant_i*. Columns (3)-(6) consider measures based on the household's financial inclusion before the reduction: *Asset Management_i* indicates if i is imputed to have participated in asset management before the reduction, based on the boosted trees imputation described in Section 5.1 applied to the SCF dataset; *Stock Market Participant_i* indicates if i is imputed to have participated in the stock market before the reduction; *Homeowner_i* indicates if i is imputed to be a homeowner; and *High Dividend Zip Code_i* indicates if the zip code associated with i 's income bracket and state of residence has an above-median share of households reporting dividend income, based on the IRS SOI Tax Stats dataset. Lower values of the financial inclusion measures proxy for greater constraints imposed by the previous account minimum. The sample consists of households who become robo participants after the reduction, except in column (2) where households who became robo participants before the reduction are also included. The remaining notes are the same as in Table 2.

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

y =	New Participant _{i,t}				
	(1)	(2)	(3)	(4)	(5)
<i>Middle_i × Post_t</i>	0.007 (0.000)	0.007 (0.000)	0.007 (0.000)		
<i>Middle_i × Monthly News Articles_t</i>		-0.000 (0.186)			
<i>Middle_i × Monthly Advisor Blogs_t</i>			-0.000 (0.101)		
<i>Middle_i × Months Before_{t,3+}</i>				0.000 (0.920)	0.000 (0.920)
<i>Middle_i × Months Before_{t,2}</i>				-0.002 (0.279)	-0.002 (0.279)
<i>Middle_i × Months After_{t,0}</i>				0.007 (0.000)	0.007 (0.000)
<i>Middle_i × Months After_{t,1}</i>				0.005 (0.015)	0.005 (0.014)
<i>Middle_i × Months After_{t,2+}</i>				0.008 (0.000)	0.008 (0.000)
<i>Millennial_i × Post_t</i>					-0.000 (0.689)
Household FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.011	0.011	0.011	0.011	0.011
Number of Observations	504,504	504,504	504,504	504,504	504,504

Note: P-values are in parentheses. This table estimates equation (7), 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

$$New\ Participant_{i,t} = \beta (Middle_i \times Post_t) + \alpha_i + \tau_t + u_{i,t},$$

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)-(3) 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 (4)-(5) 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 $Millennial_i$. The set of news articles used to construct $News\ Articles_t$ 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
<u>(a) Portfolio Parameters:</u>		
Idiosyncratic Volatility ($\sigma_{\varepsilon,i}$)	Table 7	Portfolio Dataset
Factor Loadings (β_i)	Appendix Table A4	Portfolio Dataset
<u>(b) Asset Pricing Factors:</u>		
Market Factor, Mean	0.076	CRSP
Market Factor, Volatility	0.147	CRSP
Fama-French Factors	Appendix Table A5	French
Bond Factors	Appendix Table A5	Bloomberg-Barclays
<u>(c) Preferences:</u>		
Coefficient of Relative Risk Aversion (γ)	9	Standard
Discount Factor (δ)	0.96	Standard
<u>(d) Labor Income Parameters:</u>		
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 ($\sigma_{\bar{\varepsilon}}$)	0.271	CGM
Loading on Financial Return (β^Y)	0.001	CGM
Probability of Disaster (ϕ)	0.001	SCF
<u>(e) Other Parameters:</u>		
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, summarizing 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. Panels (c) and (d) summarize preference parameters. 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 $\bar{T} = 100$; 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 10 has details on these data sources.

Table 7: Summary of Self-Managed and Robo Portfolios

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

Number of Middle-Class Portfolios: 354

Number of Upper-Class Portfolios: 1,559

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. Each panel summarizes a different statistic: 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 factor model is the Fama-French 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, and column (7) tests for a difference in this difference between wealth classes. 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 13. The remaining notes are the same as in Table 1.

Table 8: Summary of Model Fit

	Model	Data	Source
	(1)	(2)	(3)
<u>(a) Middle-Class Robo Participants:</u>			
Pre-Reduction Share	12.7%	13.4%	Deposits Dataset
Post-Reduction Share	21.3%	28.4%	
<u>(b) Robo Share:</u>			
New Middle-Class Participants	32.7%	30.6%	Deposits Dataset
New Upper-Class Participants	16.8%	16.2%	
<u>(c) Growth in Number of Participants:</u>			
All Participants	11.9%	13.4%	Table 3
Middle-Class Participants	87.1%	107.9%	

Note: This table summarizes the ability of the life cycle model from Section 6 to fit the data, which assesses the model's positive implications. Panel (a) summarizes the share of robo participants from the middle class before and after the reduction, calculated as follows: in the model, we compute the share of participants from the middle class among those who find it optimal to participate under the previous minimum and under the reduced minimum; in the data, we compute the share of participants from the middle class among those who participate before the reduction in minimum and who become participants after the reduction. Panel (b) summarizes the average portfolio share allocated to the robo advisor for 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 cumulative robo investment to liquid assets among middle and upper class households who become participants after the reduction, as in equation (6). Panel (c) 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 from the corresponding number 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. 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)
<u>(a) Baseline Model</u>			
Total Gain	1.93%	1.96%	0.01%
<u>(b) Decomposition</u>			
Priced Risk	0.29%	0.30%	0.00%
Diversification	1.26%	1.29%	0.01%
Risky Share	0.38%	0.37%	0.01%
<u>(c) Model Extensions</u>			
Participation Costs	2.33%	2.51%	0.01%
Borrowing	1.70%	1.73%	0.01%
Defined Contribution Plan	2.93%	3.25%	0.24%
<u>(d) Effect of Comparison Shocks</u>			
+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 (21). 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, defined in equation (22): the gain when moving from the pre-reduction world to a counterfactual post-reduction world in which households cannot increase their risky share and self-managed and robo portfolios only differ in their quantity of priced risk (Priced Risk); the gain when moving from the Priced Risk counterfactual to a similar counterfactual post-reduction world in which the two portfolios also differ in idiosyncratic risk (Diversification); and the gain when moving from the Diversification counterfactual to the actual post-reduction world in which households can increase their risky share (Risky Share). 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 2016 of 12% (Borrowing); and 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 under alternative shocks 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 ($v_{i,t}$) and temporary ($\zeta_{i,t}$) 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 Age

	Age 25-35		Age 36-55		Age 56-65	
	Middle	Upper	Middle	Upper	Middle	Upper
	(1)	(2)	(3)	(4)	(5)	(6)
<u>(a) Increase in Lifetime Consumption:</u>	1.81%	0.01%	1.94%	0.01%	2.23%	0.00%
<u>(b) Probability of Becoming Robo Participant by Retirement:</u>						
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 (21). 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, $P^R(\underline{T}) = 0$ 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 10 describes our data in greater detail. Appendix ?? characterizes a stylized model grounded in the framework sketched in Section ?. Appendix ?? describes how we impute variables not observed in our robo advising dataset. Appendix ?? performs an aggregated analysis referenced in Section ?. Appendix 13 describes how we estimate expected risky return as used in Section ?.

10 Data Appendix

We provide additional details on the paper’s two principal datasets: a weekly panel of deposit activity by robo participants (10.1); and a cross section of portfolio snapshots for self-managed, non-robo portfolios (10.2). We also describe other datasets (10.3.1) and provide a catalog of the paper’s key variables (10.4).

10.1 Deposits Dataset

Our 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 subjective risk tolerance, denoted *High Risk Tolerance_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.

10.2 Portfolio Dataset

The portfolio dataset contains snapshots of individuals' portfolio holdings in an outside, traditional brokerage account. This information is paired with the portfolio holdings of the individual's counterfactual automated financial management (AFM) account with our data provider, along with basic demographic information about the individual's age, annual income, and financial wealth (i.e. liquid net worth). 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, our data provider would take a snapshot of the account holdings and run an advice-generating algorithm on it. This produces a set of snapshots of individuals' non-AFM accounts. While the advice algorithm ran, our data provider would ask the individual to answer its standard questionnaire meant to gauge risk preferences, which is the source of our three demographic variables. Finally, at the conclusion of the report, our data provider would tell the individual the portfolio she would receive as a client, and give her the option to transfer. This produces a matched, counterfactual AFM portfolio for each individual in the sample. Portfolios are matched based on tax status and the client's risk preferences according to their answers to the questionnaire. Also, given this paper's research question, we are interested in this tool insofar as it provides information about non-AFM account holdings, and we do not focus on the nature of the advice received. The tool was launched in January 2016, and our sample contains 2,654 snapshots taken between January 2016 and November 2016. We merge this dataset with security level information from CRSP to produce a cross sectional dataset of individuals' brokerage and counterfactual AFM portfolios. As mentioned in the text, merging the deposits and portfolios datasets is not necessary for our analysis.

In principle, we could merge the two datasets 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.

10.3 Auxiliary Datasets

10.3.1 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 calculate two sets of variables using the SCF dataset. First, we 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 Sections ?? and ?? of the text. Table [A3](#) reports the boundaries that define the five quintiles.

The second set of variables that we calculate are measures of participation in asset management, the stock market, and homeownership. Participation in asset management is not directly observed in the SCF, and we proxy for it using the intersection of stock market participation and consulting with a broker, financial planner, banker, accountant or lawyer regarding investment. Participation in the stock market is defined as ownership of stocks, mutual funds, a trust, or an IRA. Participation in homeownership is defined as owning owner-occupied residential real estate. We calculate these variables using the SCF dataset because we do not observe them in our robo advising dataset.

10.3.2 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 13.

10.3.3 Other Auxiliary Datasets

We use the Center for Disease Control’s mortality tables to calibrate the survival probabilities 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., $t \geq T$), we use the lowest survival probability across age brackets. Lastly, we calibrate the credit card interest rate of 12% using the 2016 Federal Reserve Consumer Credit Report.

10.4 Description of 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 robo advising 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 robo advising 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 individuals 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 robo advising dataset.
- *Income_i*: This variable is annual income for household i , based on the robo advising dataset.
- *Age_i*: This variable is the age of the householder for household i , based on the robo advising 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 robo advising 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 robo advising 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 robo advising dataset.

- $\Delta Risky Share_i$: This variable is the ratio of household i 's deposits with the robo advisor over the post-reduction period (i.e., July 7, 2015 through February 29, 2016) to the household's liquid assets, based on the robo advising dataset.

11 Econometric Appendix

[Outline TBD]

11.1 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

$$g = \frac{\text{New Participants}}{\text{Existing Participants}'} \quad (\text{B1})$$

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

$$\text{New Participants} = \mathbb{E} [\text{New Participant}_i] \times \text{All Participants}, \quad (\text{B2})$$

where *All Participants* = *New Participants* + *Existing Participants* is the sum of new and existing robo participants; and $\mathbb{E} [\text{New Participant}_i]$ 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

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

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 investment constraints and, thus, $\mu = 0$. Under this counterfactual, the overall number of robo participants grows at the rate

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

or, equivalently,

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

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

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

which we can rewrite as

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

11.2 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 mismeasure $Middle_i$ as $\widehat{Middle}_i = Middle_i + \varepsilon_i$, then the estimator for μ in a specification of equation (2) without controls is

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

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

11.3 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 micro-data from each major U.S. robo advisor. The results in Appendix Figure 13 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.

11.4 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.

12 Theory Appendix

[Outline TBD]

12.1 Model Solution

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

12.1.1 Consolidated Problem

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

$$V_t(W_{i,t}) = \max_{\alpha_{i,t}^f, \alpha_{i,t}^S, \alpha_{i,t}^A} \left\{ \frac{C_{i,t}^{1-\gamma}}{1-\gamma} + \delta p_t \mathbb{E}_t [V_{t+1}(W_{i,t+1})] \right\} \quad (\text{C1})$$

s.t. (9)-(17)

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

$$V_{\underline{T}+1}(W_{i,\underline{T}+1}) = \max_{\{\alpha_\tau^f\}} \sum_{\tau=\underline{T}+1}^{\tau=\bar{T}} (\delta p_{\underline{T}+1})^{\tau-\underline{T}-1} \frac{C_{i,\tau}^{1-\gamma}}{1-\gamma} \quad (\text{C2})$$

s.t.

$$0 \leq \alpha_\tau^f \leq 1 \quad (\text{C3})$$

$$W_{i,\tau+1} = \alpha_{i,\tau}^f (1 + R^f) W_{i,\tau} \quad (\text{C4})$$

$$C_{i,\tau} = [1 - \alpha_{i,\tau}^f] W_{i,\tau}. \quad (\text{C5})$$

Indirect utility has the familiar form

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

with

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

$$\chi = \delta^{\frac{1}{\gamma}} (1+R^f)^{\frac{1-\gamma}{\gamma}} \quad (\text{C8})$$

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

12.1.2 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 90th 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. Likewise, we discretize the set of shocks and approximate their joint distribution through Gaussian quadrature (e.g., [Tauchen and Hussey \(1991\)](#)). For completeness, the model's shocks are: $F_t, \epsilon_{i,t}^S, \epsilon_{i,t}^A$, and $\bar{v}_{i,t} \equiv \bar{\zeta}_{i,t} + v_{i,t}$.

Following standard practice, we optimize by grid search, and so we avoid selecting local optima. Accordingly, we discretize the control variables: $\alpha_{i,t}^f, \alpha_{i,t}^S$, and $\alpha_{i,t}^A$. The control variable grid omits choices that violate one of the constraints (9)-(17). As mentioned in the text, we simplify the model's computational complexity by assuming households cannot hold the self-managed and automated portfolios concurrently: $\min \{ \alpha_{i,t}^S, \alpha_{i,t}^A \} = 0$. We obtain similar results without this simplification because it is rarely optimal to hold both at the same time, given their empirical covariance. We also reduce simplify computational complexity by assuming households must maintain a minimum balance of M with the automated asset manager, as distinct from a minimum investment. Otherwise, we must keep track of α_t^A as an additional state variable because it determines the lower bound on a household's robo investment. Namely, a pure minimum investment requirement does not require households 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 pure minimum investment 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 $V_T(W_{i,T})$, we iterate backward, solving $V_{T-1}(W_{i,T-1})$ 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 $\{ \alpha_{i,\tau}^f(W_{i,\tau}), \alpha_{i,\tau}^S(W_{i,\tau}), \alpha_{i,\tau}^A(W_{i,\tau}) \}$ that we use in the positive and welfare analyses of [Sections 7 and 8](#).

12.2 Welfare Measure

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

$$q_i = \left(\frac{\bar{V}_i}{\underline{V}_i} \right)^{\frac{1}{1-\gamma}} - 1, \quad (\text{C9})$$

as shown in [Appendix 12](#).

Like other papers that use life cycle models to study distributional effects (e.g., [Gete and Zecchetto \(2018\)](#)), we recalculate equation (C9) separately for subpopulations defined by wealth and age. Then, we decompose equation (C9) into three terms that reflect the particular gain from changes in diversification,

priced risk, and risky share. Explicitly,

$$q_i = \underbrace{\left[\left(\frac{\bar{V}_i |_{\alpha, \sigma_{\epsilon,i}^P}}{V_i} \right)^{\frac{1}{1-\gamma}} - 1 \right]}_{\text{PricedRisk}} + \underbrace{\left[\left(\frac{\bar{V}_i |_{\alpha}}{V_i} \right)^{\frac{1}{1-\gamma}} - \left(\frac{\bar{V}_i |_{\alpha, \sigma_{\epsilon,i}^P}}{V_i} \right)^{\frac{1}{1-\gamma}} \right]}_{\text{Diversification}} + \underbrace{\left[\left(\frac{\bar{V}_i}{V_i} \right)^{\frac{1}{1-\gamma}} - \left(\frac{\bar{V}_i |_{\alpha}}{V_i} \right)^{\frac{1}{1-\gamma}} \right]}_{\text{RiskyShare}} \quad (\text{C10})$$

where $\alpha_i \equiv \alpha_i^S + \alpha_i^A$ is household i 's risky share; $\bar{V}_i |_{\alpha}$ 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, $\bar{V}_i |_{\alpha, \sigma_{\epsilon,i}^P}$ 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 (C10) 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., $\beta_i^A \neq \beta_i^S$). The second term equals the gain when the two portfolios also differ in idiosyncratic risk (i.e., $\sigma_{\epsilon,i}^A \neq \sigma_{\epsilon,i}^S$), such that portfolio returns equal their actual values according to equation (9), but households still cannot increase their risky share. Notably, these first two terms equals zero for households who do not already 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.

12.3 Model Extensions

Section 8.1.3 considers various model extensions that we now describe.

12.3.1 Participation Costs

The Participation Cost extension introduces a per-period utility loss from self-management equal to a \$100 reduction in consumption.

12.3.2 Borrowing

The Borrowing extension allows borrowing up to 30% of liquid wealth at a rate of 12% (c.f. FRED), which must be repaid the following period (e.g., lender-friendly bankruptcy laws allow recourse up to 30% of liquid wealth).

12.3.3 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. This setup and parameterization follows Campbell et al (2001). Liquid wealth is defined to include the value of plan assets accumulated to that point. The plan cannot be liquidated.

12.4 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 ten moments we seek to match are ...

13 Asset Pricing Appendix

This appendix describes the method for estimating the expected return on robo portfolios in Section ???. As mentioned in Section ??, using a historical average to measure expected return is subject to well-known issues related to limited time horizons, and especially so in our setting given the relatively-short history of many ETFs managed by our data provider. Therefore, we follow [Calvet, Campbell and Sordini \(2007\)](#) and propose an asset pricing model to estimate the expected returns for the securities in our sample. While imposing a model improves the efficiency of expected return estimates relative to directly measuring them from historical returns, it leads to some bias by imposing an imperfect model of the return structure. Since the choice of model is somewhat arbitrary and the degree of bias will depend on the characteristics of the portfolio in question, we estimate expected return across a variety of common asset pricing models, indexed by their vector of factors F . For each security k (i.e., ETF k), we estimate the following pricing kernel

$$Return_{k,t} = \beta_k^F F_t + \epsilon_{k,t}^F \tag{D1}$$

where F_t denotes a column vector of pricing factors in month t ; β_k^F denotes the respective row vector of loadings; $Return_{k,t}$ 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 $\epsilon_{k,t}^F$ is an idiosyncratic, zero-mean shock to security k with standard deviation σ_k^F . We estimate equation (D1) using the longest available time series of monthly returns for each security k and factor vector F dating back to January 1975. Given the estimated loadings $\hat{\beta}_k^F$ from estimating equation (D1) for model F , it is straightforward to compute the expected return on household i 's robo portfolio, denoted $Risky\ Return_i^F$. Explicitly, if there are K securities and N factors, then

$$Risky\ Return_i^F = w_i \beta^F \pi^F, \tag{D2}$$

where w_i is a $1 \times K$ row vector of weights across securities in household i 's robo portfolio; β^F is a $K \times N$ matrix of factor loadings; and π^F is an $N \times 1$ column vector of factor risk prices. We estimate equation (D1) for the following asset pricing models,

$$F_t^{CAPM} = [R_t^m]', \tag{D3}$$

$$F_t^{FF} = [R_t^m, R_t^{HML}, R_t^{SMB}]', \tag{D4}$$

$$F_t^{FF+} = [R_t^m, R_t^{HML}, R_t^{SMB}, R_t^{USB}, R_t^{GLB}]', \tag{D5}$$

where R_t^m is the monthly market return based on the global Morgan Stanley Capital International Index (MSCII), net of the one-month Treasury yield; R_t^{HML} is the spread in monthly return between high book-to-market stocks and low book-to-market stocks; R_t^{SMB} denotes the spread in monthly return between stocks with a small market capitalization and a big market capitalization; R_t^{USB} is the monthly return on the Barclays Aggregate U.S. Bond Index Unhedged, net of the one-month Treasury yield; and R_t^{GLB} is the monthly return on the Barclays Global Aggregate Bond Index Unhedged, net of the one-month Treasury yield. In words, equations (D3)-(D5) 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 10. We use the sample mean to calibrate the factor risk prices, multiplied by 12 to obtain an approximate annual value. The corresponding risk prices

for the models in equations (D5)-(D5) are

$$\pi^{CAPM} = [0.068]', \quad (D6)$$

$$\pi^{FF} = [0.068, 0.036, 0.004]', \quad (D7)$$

$$\pi^{FF+} = [0.068, 0.036, 0.004, 0.062, 0.060]'. \quad (D8)$$

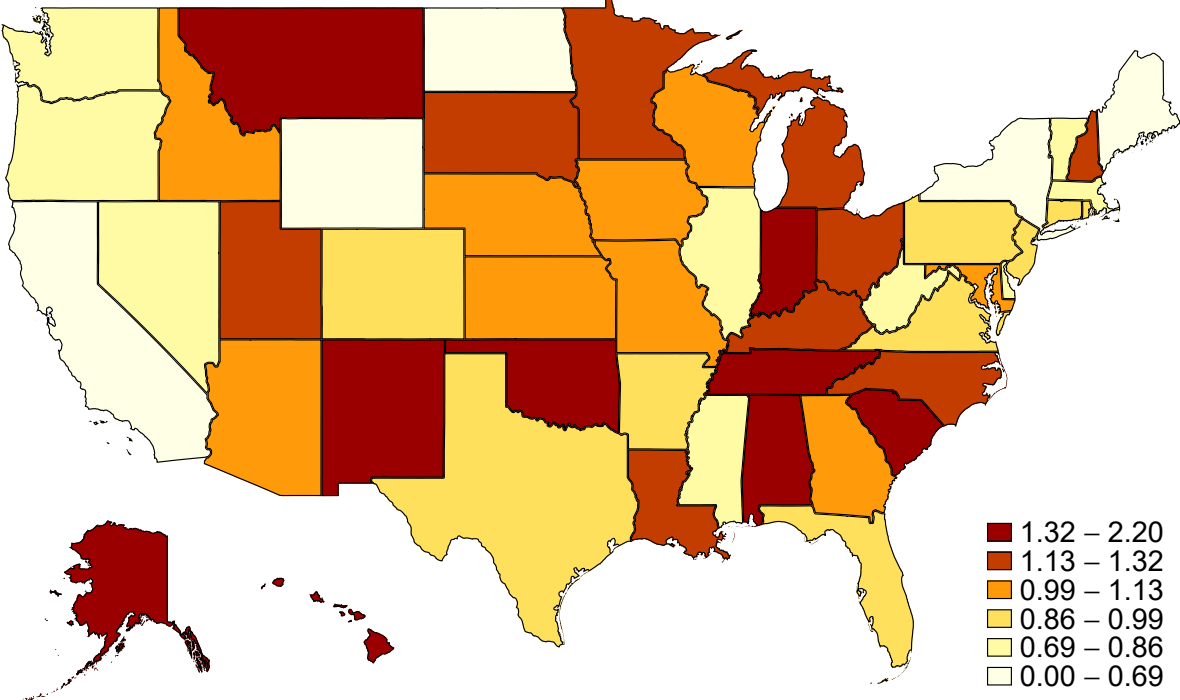
Similarly, we use the sample covariance matrix to calculate the share of risk that is compensated according to a given asset pricing model, studied in Section ???. Explicitly, this share equals

$$\text{Compensated Risk Share}_i^F = \frac{(w_i \beta^F) \Sigma^F (w_i \beta^F)'}{(w_i \beta^F) \Sigma^F (w_i \beta^F)' + w_i \Sigma_\epsilon^F w_i'} \quad (D9)$$

where Σ^F is an $N \times N$ covariance matrix of factor returns under asset pricing model F ; and Σ_ϵ^F is an analogous $K \times K$ covariance matrix of idiosyncratic returns, $\epsilon_{k,t}^F$.

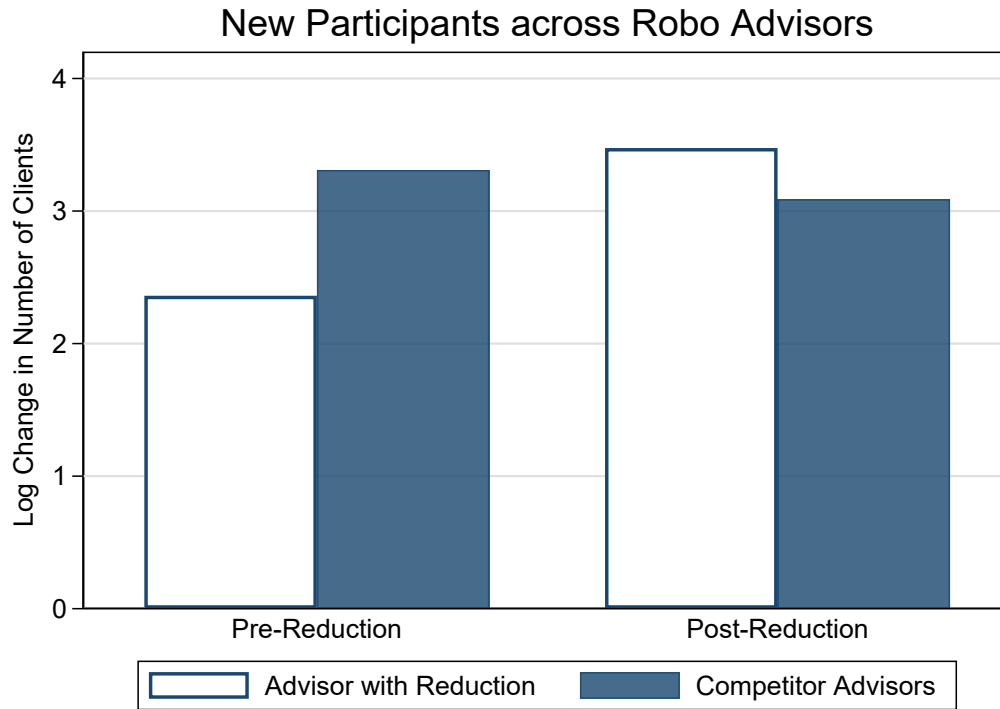
Additional Figures and Tables

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



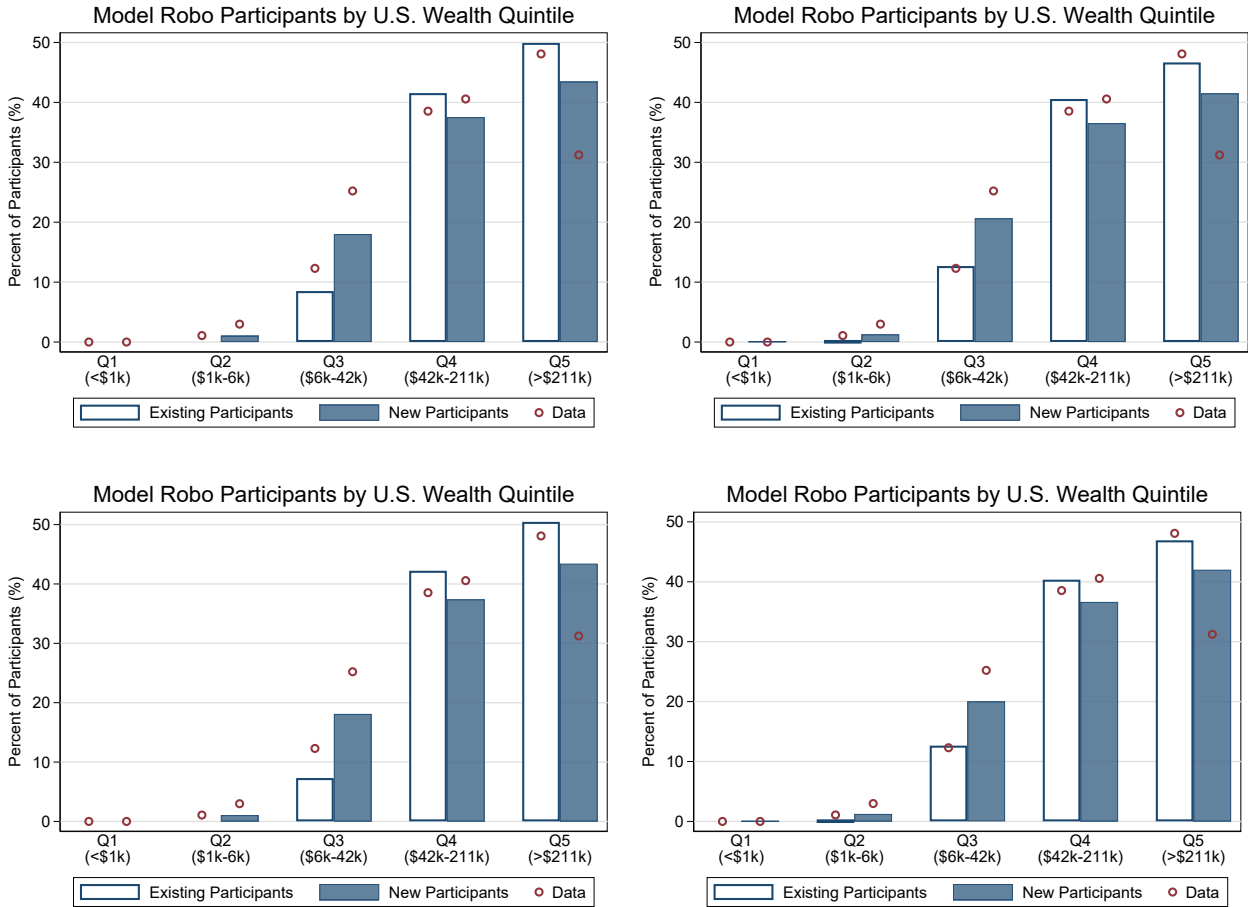
Note: This figure plots the change in the log of the number of robo participants from each U.S. state plus one. The change is from the pre-reduction period (December 1, 2014 to July 7, 2015) to the post-reduction period (July 7, 2015 to February 29, 2016).

Figure A2: Reallocation of Participants across Robo Advisors



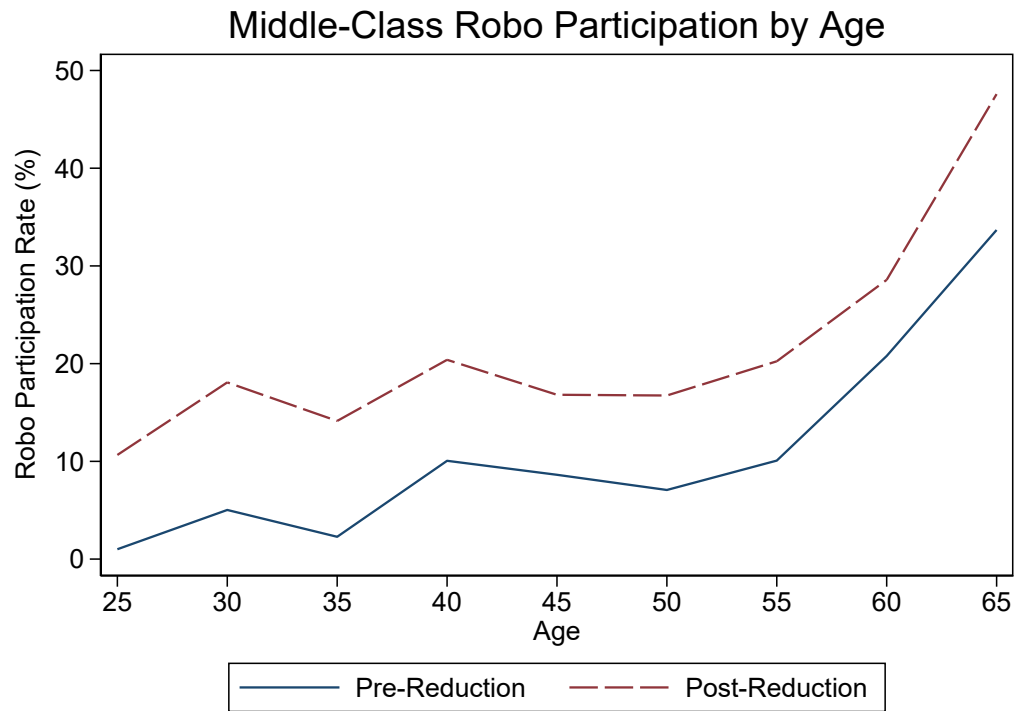
Note: This figure plots the log of the change in the number of clients across robo advisors, in thousands, which assesses whether the reduction increases robo participation or simply reallocates robo participants across advisors. The change is calculated separately for the robo advisor that reduced its 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 Vector



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 models. The factor model in panels (a)-(b) is the CAPM, calibrated using the 1960-2017 period in panel (a) and the 2010-2017 period in panel (b). The factor model in panel (c) is the Fama-French Three Factor Model (Fama-French). The factor model in panel (d) is the baseline Fama-French model augmented with U.S. and global bond returns, and so panel (b) shows the same figure as in Figure 5. The moments of the asset pricing factors are summarized in Appendix Table A5. The remaining notes are the same as in Figure 5.

Figure A4: Optimal Middle-Class Robo Participation 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 minimum investment (\$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. 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 minimum investment 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 13. Column (6) shows the percent of robo participants with the indicated portfolio. Column (7) shows the participants's age. 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 in Asset Management (%)	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 professional asset management, 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: Summary of Non-Robo and Robo Portfolio Loadings

Factor Loadings (β_i)	Middle Class			Upper Class		
	Non-Robo	Matched Robo	Difference	Non-Robo	Non-Robo	Difference
Market	0.930	0.893	-0.036 (0.002)	0.899	0.876	-0.023 (0.000)
SMB	0.044	0.003	-0.040 (0.004)	0.030	-0.001	-0.032 (0.000)
HML	-0.086	0.061	0.147 (0.000)	-0.074	0.061	0.135 (0.000)
GLB	0.629	-0.020	-0.649 (0.000)	0.574	-0.023	-0.597 (0.000)
USB	-0.447	0.508	0.955 (0.000)	-0.391	0.512	0.903 (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 A5: 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. Panel (a) summarizes the covariance matrix: 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, based on data from Ken French's website; SMB is the spread in monthly return between stocks with a small market capitalization and a big market capitalization, based on data from Ken French's website; 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. 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%.

Table A6: Summary of Self-Managed and Robo Portfolios by Asset Pricing Factor

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

Number of Middle-Class Portfolios: 354

Number of Upper-Class Portfolios: 1,559

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. Each panel summarizes a different statistic: 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 factor models are: the CAPM and the Fama-French Three Factor Model (Fama-French). 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, and column (7) tests for a difference in this difference between wealth classes. 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 13. The remaining notes are the same as in Table 1.

Table A7: Robustness of Welfare Implications

	Increase in Lifetime Consumption (%)		
	All Households	Middle Class	Upper Class
	(1)	(2)	(3)
<u>Parameterization</u>			
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 A8: GMM Estimates of Preference Parameters

Parameter	Estimate	
	(1)	(2)
Coefficient of Relative Risk Aversion (γ)	9.1 [8.868, 9.332]	8.5 [8.242, 8.758]
Discount Factor (δ)	0.945 [0.928, 0.962]	0.990 [0.970, 1.000]
Weights	GMM Optimal	Equal

Note: This table estimates the life cycle model's preference parameters using the generalized method of moments estimator described in Appendix 12. 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. Standard errors are bootstrapped. Brackets correspond to 95% confidence intervals under an asymptotically normal distribution.