



Robo advisors and access to wealth management[☆]

Michael Reher^{a,*}, Stanislav Sokolinski^b

^a University of California San Diego, Rady School of Management, United States of America

^b Broad College of Business, Michigan State University, United States of America

ARTICLE INFO

Dataset link: <https://data.mendeley.com/datasets/3f88g2cb7x/2>

JEL classification:

G11

G24

D3

O3

Keywords:

FinTech

Financial advice

Portfolio delegation

Inequality

ABSTRACT

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

“The wealth-management industry stratifies customers in a manner rather similar to airlines. ‘High-net-worth’ clients fly business class, picking stocks and chatting in person with named advisors. Cattle class gets no service at all. Technology is conspiring to change that.”

The Economist Magazine

1. Introduction

Recent technological innovation has enabled financial intermediaries to scale the provision of many traditional services. In the wealth

management market, this scaling has taken the form of “robo advisors”: intermediaries that use automation to provide services similar to a traditional wealth manager (D’Acunto and Rossi 2020). Robo advisors have grown rapidly in both popularity and size over the past decade. For example, the largest five robo advisors have grown roughly tenfold, and all of the Big-4 U.S. banks now offer a robo advising service.¹

Robo advisors claim to have two key advantages over their traditional counterparts. First, by adhering to a transparent investing algorithm, robo advisors can benefit their clients by circumventing inefficiencies documented among financial advisors, many of which stem

[☆] Nikolai Roussanov was the editor for this article. We thank Christoph Aymanns, John Campbell, Michael Boutros, Francesco D’Acunto, Ben Friedman, Pedro Gete, Steffen Hitzemann, Jakub Jurek, Iva Kalcheva, Yaron Levi, Jun Liu, Matteo Maggiori, Apurva Mehta, Will Mullins, Alessandro Previtero, Franklin Qian, Enrichetta Ravina, Jonathan Reuter, Alberto Rossi, Nikolai Roussanov, Andrei Shleifer, Celine Sun, Allan Timmermann, Ross Valkanov, Boris Vallee, Ken Wilbur, anonymous referees, and seminar participants at UC Irvine Finance Conference, Michigan State University, University of Houston, Indiana University, University of Texas at Dallas, CEAR-RSI Household Finance Workshop, Northeastern University Finance Conference, University of Mannheim, USC, University of Wisconsin, Jackson Hole Winter Finance Conference, FINRA NORC Diversity Conference, FMA Annual Conference, Paris FinTech Conference, Toronto FinTech Conference, the New York Fed’s FinTech Conference, Georgetown’s Finance and Policy Center, UC San Diego, Rutgers Business School, California Corporate Finance Conference, CAFR FinTech Workshop, and Harvard’s finance PhD lunch for comments. Ruikai Ji provided excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the position of Wealthfront Inc.

* Corresponding author.

E-mail addresses: mreher@ucsd.edu (M. Reher), sokolins@msu.edu (S. Sokolinski).

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

from portfolio-by-portfolio discretion.² Second, because automation lowers per-portfolio management costs, robo advisors can extend the benefits of professional wealth management to less-wealthy investors, who do not have enough assets to invest with traditional managers. Our paper examines the latter claim, which is popular and intuitive but has not been studied. We use a quasi-experiment and a quantitative model to trace out the effect of access to robo advisors down to the welfare of new, less-wealthy robo investors.

In the first part of the paper, we study how access to robo advisors affects participation by the less-wealthy. Not only is this exercise important in its own right, but the results will also help discipline the model that we develop in the second part of the paper. However, identifying the effect of interest is challenging for two reasons: lack of standardized data on robo participants; and a need for a setting in which robo advisors suddenly become more accessible to the less-wealthy. We overcome these challenges by obtaining a novel dataset directly from a major U.S. robo advisor and by studying a quasi-experiment. In our setting, described in Sections 2 and 3, the same robo advisor suddenly reduces its account minimum from \$5,000 to \$500. This shock represents one of the first examples in which sophisticated wealth management becomes available to a wide range of less-wealthy investors.

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

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

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

In the second part of the paper, we develop and calibrate a quantitative model of asset allocation with endogenous portfolio delegation and an account minimum. This exercise addresses two major questions that cannot be answered by the reduced-form DiD analysis from the

first part. First, we ask whether differences in asset allocation suffice to explain the increase in middle-class participation with the advisor, or whether we must appeal to non-financial channels, like peace-of-mind from portfolio delegation (e.g., Gennaioli et al. 2015). Second, we assess the welfare gain from access to robo advisors, using the standard lifetime consumption metric (e.g., Gomes 2020). This approach lets us examine the potential benefits of robo advisors more comprehensively than, say, using realized returns, as it accounts for the investor's horizon and human capital.³ We are particularly interested in how these potential benefits may vary across investors and which features of robo portfolios drive the results.

Our comparison group consists of retail investors who would otherwise manage their risky assets on their own without access to a robo advisor (“self-managed portfolios”). Section 6 uses portfolio-level microdata to document three advantages that robo portfolios have over self-managed ones. First, robo portfolios contain greater exposure to priced risk factors, such as bond and value premia. Consequently, they provide a 2 pps higher expected return. Second, they are much better diversified. Third, robo portfolios exhibit personalization by both age and wealth, which we call a “double glide path”: exposure to stocks falls as an investor grows older, holding wealth fixed; but it rises as an investor becomes wealthier, holding age fixed. So, robo portfolios are more personalized than target date funds (TDF), which offer a single glide path by age.

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

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

² A partial list of inefficiencies includes impulsive recommendations (e.g., Linnaïmaa et al. 2021), pandering to client biases (e.g., Mullainathan et al. 2012), overly generic recommendations (e.g., Foerster et al. 2017), and recommendations of inferior products due to commissions (e.g., Egan 2019; Chalmers and Reuter 2020). D'Acunto and Rossi (2020) discuss how robo advisors can correct some of these inefficiencies.

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

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

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

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

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

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

Related literature

This article speaks to a nascent literature on FinTech intermediaries, which includes robo advisors. Unlike prior and contemporaneous

studies, we study automated wealth management in a pure form that involves: complete portfolio delegation, as opposed to non-binding advice (e.g., D’Acunzio et al. 2019; Bianchi and Bri  re 2022; D’Hondt et al. 2020); no option for human interaction (e.g., Rossi and Utkus 2021b); and robo advisors unaffiliated with the banking system (e.g., Loos et al. 2020). We also focus on less-wealthy investors in a quasi-experimental setting, complementing the literature’s typical focus on more-affluent investors.⁶

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

Our study also delivers a new perspective on the drivers of demand for professional portfolio management. The literature has argued that retail investors seek portfolio managers because managers actively collect information about the underlying assets (e.g., G  rleanu and Pedersen 2018) or confer peace of mind (e.g., Gennaioli et al. 2015). We find that many less-wealthy investors have more mundane needs: to construct personalized and diversified portfolios, even if the underlying assets are passively managed.⁸ These mundane needs are not fully met by another form of personalized asset management, TDFs (e.g., Balduzzi and Reuter 2019; Parker et al. 2021), as robo advisors offer more personalization via a double glide path.

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

⁶ Our study also relates to a broader question of how technological innovation affects financial inclusion and wealth inequality. See, for example, the studies in the contexts of app-based payments (e.g., Hong et al. 2022), bank deposits (e.g., Bachas et al. 2018; Bachas et al. 2020; Higgins 2022), and mortgage markets (e.g., Fuster et al. 2019; Bartlett et al. 2021; Fuster et al. 2021).

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

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

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

like: preferences (e.g., Barberis et al. 2006), sophistication (e.g., Grinblatt et al. 2011; Christelis et al. 2010), socialization (e.g., Hong et al. 2004), or education (e.g., Cole et al. 2014; Van Rooij et al. 2011).

2. Institutional background

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

2.1. The U.S. robo advising market

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

Appendix Table A1 summarizes the largest robo advisors in the U.S. as of July 2015, including their account minimums, assets under management, fees, and provision of traditional, human-based management. At the time of our analysis, Wealthfront managed \$2.4 billion and was the largest standalone robo advisor in the U.S. market, with Betterment and Personal Capital as its nearest competitors. Wealthfront is also the only robo advisor that relies purely on automation, with no option for a human advisor. Two traditional managers, Vanguard and Charles Schwab, launched robo advising services early in 2015. Both of these services managed more than Wealthfront because they transferred assets from existing, non-robo services. From here on, we refer to Wealthfront simply as "the robo advisor".

2.2. The robo advisor's baseline service

We study the robo advisor's baseline service: an automatically rebalanced portfolio of 10 ETFs corresponding to 10 asset classes. The 10 ETFs are chosen to track stock market indices (VIG, VTI, VEA, VWO), bond market indices (LQD, PCY, MUB, SCHP), and other asset classes, namely real estate (VNQ) and commodities (XLE). The portfolio weights are determined by a questionnaire about the client's age, liquid assets, income, demographic background, and response to hypothetical investment decisions. The client is then assigned to one of 20 possible risk tolerance scores, which range from 0.5 to 10 in increments of 0.5. Each risk tolerance score uniquely determines a robo portfolio. The portfolio weights solve a problem of optimal asset allocation across the 10 ETFs, taking this score as a parameter.¹⁰

⁹ The robo advisor we study is an automated wealth manager because it offers an entire pipeline of wealth management services ("wealth manager") without human interaction ("automated"), as opposed to just providing financial advice. This pipeline includes the baseline product that we study (a personalized and automatically rebalanced portfolio) as well as: tax loss harvesting; a risk parity fund; and financial planning on both the asset and liability sides of the investor's balance sheet.

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

Importantly, robo portfolios are not recommendations: rather, they are directly selected, managed and rebalanced by the robo advisor. Consequently, investors have little discretion over their portfolio allocations, and so their robo performance will not depend on sophistication (e.g., Grinblatt et al. 2011; Christelis et al. 2010), ability to diversify (e.g., Calvet et al. 2007), willingness to follow advice (e.g., Bhattacharya et al. 2012), or reluctance to rebalance (e.g., Calvet et al. 2009). For its services, the advisor charges an asset-based management fee of 0% for accounts under \$10,000 and 0.25% for larger accounts. So, most of the less-wealthy investors in our data incur no management fee.

2.2.1. Risk profile of robo portfolios

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

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

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

2.3. The 2015 reduction in account minimum

On July 7, 2015, the robo advisor unexpectedly reduced its account minimum from \$5,000 to \$500, meaning that an investor would need to invest only \$500 to participate with the advisor as opposed to \$5,000 beforehand. This reduction is sizeable from the standpoint of most U.S. retail investors. For reference, \$5,000 equals 30% of median liquid assets (\$17,000), and it defines the 37th percentile of the U.S. wealth distribution, according to the 2016 Survey of Consumer Finances. So, prior to the reduction, half of U.S. retail investors could not participate without investing at least 30% of their wealth, while 37% could not participate at all without borrowing.

At the time of the reduction, all of the largest five U.S. robo advisors required an account minimum of at least \$5,000 except for one, which had no account minimum but maintained a fee structure that discouraged setting up small accounts.¹¹ Importantly, the month of the reduction does not coincide with any other product launches by the robo advisor, any changes in its fee, or any significant developments in

return with the primary ETF. The correlation coefficient is quite high and often greater than 0.99.

¹¹ Betterment charged a \$3 service fee on accounts under \$10,000 for customers who do not auto-invest \$100 monthly in their accounts. This fee structure implies a 7.2% annual management fee for a \$500 account and a 36% management fee for a \$100 account.




Fig. 1. Understanding the robo advisor's asset allocation.

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

3. Data

Our analysis relies on two core datasets that draw from the same population: a panel dataset covering deposit activity by investors who participate with the robo advisor (3.1); and a dataset on self-managed, outside portfolio holdings that includes both investors in the deposits

dataset and robo non-participants (3.2). Appendix A has details, including information on additional datasets that we use. We use the term "robo participant" to describe investors who have invested money with the robo advisor.

3.1. Deposits dataset

The first core dataset contains a weekly time series of deposits with the robo advisor from December 1, 2014 through February 29, 2016. Straddling the July 2015 reduction, this time window also marks a formative period in the history of the robo advising market when the number of robo participants was still small. For each investor, we observe the size of their weekly deposits and the following demographic variables: annual income; state of residence; age; and liquid assets, defined as "cash, savings accounts, certificates of deposit, mutual funds,

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

IRAs, 401ks, and public stocks". The demographic variables are self-reported via the robo advisor's questionnaire and are static. We address measurement error from potential misreporting later in Section 4.3. To match the setup of our model, in which account holders have at most one risky account, we treat taxable and non-taxable accounts of the same account holder in our deposits dataset as separate observations (separate "investors"). Section 5.3 reproduces our main results after dropping the 12% of account holders who have multiple accounts.

Studying a private, company-specific dataset has two major advantages over publicly available datasets such as the SEC's Form ADV filings, which serve as the basis for many industry reports about the robo market. First, Form ADV data includes inactive clients who stopped making deposits as well as "clients" who create a username but never provide the robo advisor with any capital to manage.¹³ Second, our dataset includes information about a robo participant's wealth, allowing us to study investment activity across the wealth distribution.

We use the 2016 Survey of Consumer Finances (SCF) to categorize investors into three wealth groups. We respectively use the terms "lower class", "middle class", and "upper class" to describe investors from the first, second or third, and fourth or fifth quintiles of the U.S. distribution of liquid assets. The resulting boundary between the lower versus middle class is \$1,000, and the boundary between the middle versus upper class is \$42,000.¹⁴ Since the SCF includes a representative sample of U.S. retail investors, this approach ensures that we define wealth groups in accordance with the overall U.S. population.

3.1.1. Summary statistics

Table 1 summarizes the deposits dataset. Panel (a) compares new participants who join the robo advisor after the reduction in account minimum with existing participants. New participants are significantly less wealthy, earn lower incomes, make smaller initial deposits, and are 16 pps more likely to belong to the middle class. They are also more likely to come from less financially developed U.S. states (Appendix Fig. A1). Importantly, new robo participants invest with persistence, as 98% of them do not close their account over our sample period; and 71% make subsequent deposits that are typically much smaller than their initial investment. These results suggest that new participants did not join the advisor for a short-term experimentation.


Panel (b) conveys similar patterns within the middle-class, with two principal findings. First, over half of existing middle-class participants invested exactly \$5,000 for their initial deposit. This "bunching" suggests that the previous minimum imposed a binding constraint on their investment behavior. Second, new middle-class participants are not significantly younger than existing ones, suggesting the reduction does not disproportionately affect the young. We explore both findings later, in our quantitative model.

3.2. Portfolio dataset

The second core dataset contains information on self-managed portfolio holdings for a subset of investors who use the advisor's online financial advice tool. The tool works by loading data from the investor's brokerage account and providing free financial advice about the portfolio in that account. The tool also shows the investor the portfolio they would receive if they became a client of the advisor. Some investors decide to join the robo advisor, migrating into the deposits dataset. Other

¹³ For example, we observe 9,702 participants in our dataset, in contrast to the 61,000 reported in publicly available SEC filings. This discrepancy reflects how: "The definition of 'client' for Form ADV states that advisors must count clients who do not compensate the advisor" (SEC 2017).

¹⁴ See Appendix Table A4 for detail. Liquid assets in the SCF are calculated to match the definition in our robo advising dataset as closely as possible: the sum of checking accounts, savings accounts, certificates of deposit, cash, stocks, bonds, savings bonds, mutual funds, annuities, trusts, IRAs, and employer-provided retirement plans.



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

Fig. 2. How the robo wealth distribution changes after expanding access.

investors decide not to become robo participants. The combined dataset comprises 1,913 pairs of self-managed and counterfactual robo portfolios for both robo participants and non-participants. Investors in the portfolio dataset have similar demographic characteristics as those in the deposit dataset (Appendix Table A5).

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

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

4. Effect of accessibility on robo participation

4.1. Shift in wealth distribution of robo participants


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

Fig. 3 compares the distributions of new and existing participants across U.S. wealth quintiles. The effect of the reduction is concentrated among the middle-class. The share of robo participants from the second and third quintiles of the U.S. wealth distribution grows by 107%

Table 1
Summary of robo participants.

	Existing participants ($N = 4,366$)			New participants ($N = 5,336$)			New vs. existing	
	Mean	Standard Deviation	Median	Mean	Standard Deviation	Median	Difference in Mean	Standard Error
(a) All Investors:								
<i>Liquid Assets_i</i> ('000)	436.44	660.82	200	265.21	480.25	100	-171.22***	(11.968)
<i>Income_i</i> ('000)	157.36	110.67	130	116.17	95.9	90	-41.18***	(2.128)
<i>Initial Deposit_i</i> ('000)	33.68	94.54	10	22.56	72.61	5	-11.12***	(2.424)
<i>Age_i</i>	35.79	8.72	34	35.4	9.97	33	-0.39***	(0.190)
<i>Middle_i</i>	0.15	0.35		0.3	0.46		0.156***	(0.008)
<i>First Deposit at \$5k_i</i>	0.33	0.47		0.11	0.32		-0.212***	(0.012)
<i>First Deposit < \$5k_i</i>				0.42	0.49			
<i>No Account Closure_i</i>				0.98	0.15			
<i>Subsequent Deposit_i</i>				0.71	0.45			
(b) Middle Class:								
<i>Liquid Assets_i</i> ('000)	23.23	11.68	25	19.71	11.36	18	-3.527***	(0.542)
<i>Income_i</i> ('000)	92.86	62.21	80	67.14	42.52	60	-25.720***	(2.682)
<i>Initial Deposit_i</i> ('000)	7.6	5.34	5	4.95	12.58	2	-2.652***	(0.430)
<i>Age_i</i>	30.33	6.33	29	30.04	7.07	28	-0.293	(0.306)
<i>First Deposit at \$5k_i</i>	0.55	0.5		0.09	0.29		-0.460***	(0.028)
<i>First Deposit < \$5k_i</i>				0.66	0.47			
<i>No Account Closure_i</i>				0.97	0.18			
<i>Subsequent Deposit_i</i>				0.72	0.45			

Note: This table summarizes investors who participated with the robo advisor before the reduction in account minimum (Existing Participants) and who become participants after the reduction (New Participants), based on the Deposits Dataset. Subscript i indexes investor. *Liquid Assets_i* is the sum of cash, savings accounts, certificates of deposit, mutual funds, IRAs, 401ks, and public stocks, in thousands of dollars. *Income_i* is annual investor income, in thousands of dollars. *Initial Deposit_i* is the value of the investor's initial deposit, in thousands of dollars. *Age_i* is the investor's age. *High Risk Tolerance_i* indicates if the investor chooses a higher risk tolerance score than that recommended by the robo advisor. *Middle_i* indicates if i belongs to the second (\$1k-\$6k) or third U.S. wealth quintile (\$6k-\$42k). *No Account Closure_i* indicates if the investor does not close the account, over the sample period. *Subsequent Deposit_i* indicates if the investor makes a subsequent deposit after her initial one. Most (86%) of subsequent deposits are less than the initial deposit. *First Deposit at \$5k_i* indicates if the investor's initial deposit equals \$5,000 or is no more than 5% higher. *First Deposit < \$5k_i* indicates if the investor's initial deposit lies between \$500 and \$5,000. Wealth consists of liquid assets, and wealth quintiles are calculated using the 2016 Survey of Consumer Finances (SCF) dataset. The sample consists of investors who participate with the robo advisor and make a deposit over the period from December 2014 through February 2016. The upper panel summarizes all investors in the sample, and the lower panel summarizes investors from the second or third U.S. wealth quintile. Appendix A has details.



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

Fig. 3. Change in representativeness of robo wealth distribution.

(16 pps), while the share from the upper two quintiles falls by 18% (16 pps). This distributional shift reflects an increase in the number of

middle-class participants, not a decrease in the number of upper-class participants (Appendix Fig. A2). The democratization is incomplete, though, since investors from the first quintile remain non-participants. We find similar results using the subsample of participants in the portfolio dataset (Appendix Fig. A3).

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

4.2. Difference-in-difference setup

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

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

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

Our parameter of interest is μ , which equals the double difference in the probability of becoming a robo participant after versus before

the reduction between middle-class versus upper-class investors. We hypothesize that the reduction affects robo participation by relaxing minimum-account constraints. Accordingly, we interpret middle-class investors as the treatment group and upper-class investors as the control group.

The additional terms in equation (1) capture channels distinct from minimum-account constraints. The investor fixed effect ζ_i captures slow-moving characteristics that predispose investors to participating with the advisor, such as sophistication or trust. The interaction $X_i \times Post_t$ captures heterogeneous trends by observed investor characteristics. If, for example, younger investors are more likely to become robo participants after the reduction for reasons apart from a relaxation of minimum-account constraints, then the coefficient ψ would capture this effect.

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

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

where $New Participant_i$ indicates if investor i becomes a robo participant after the reduction; and $u_i \equiv \Delta v_i$. The following identification assumption allows us to interpret μ as the causal effect of the reduction:

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

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

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

4.3. First-difference results

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

In column (4) we estimate equation (2) using the portfolio dataset. In this case, we need not condition the interpretation of μ on eventual participation because we observe non-participants as well. We find the reduction raises middle-class investors' probability of participation by 10.6 pps in this sample. The slightly smaller effect suggests that middle-

class investors who use the online advice tool may have weaker demand for robo advisors.¹⁵

Our treatment exposure variable, $Middle_i$, may be subject to additive measurement error due to self-reporting. As we show formally in Appendix B.2, such measurement error can bias the estimate upwards only if new participants underreport their wealth relative to existing participants. We mitigate this concern by remeasuring $Middle_i$ in two ways. First, we redefine the middle class exclusively as the second quintile of the U.S. wealth distribution and omit investors from the third quintile from the sample. Under this definition, upper-class investors would need to underreport liquid assets by at least \$36,000 to be misclassified as middle-class. Second, we exclude investors whose liquid assets are within a 10% buffer of the boundary between the third and fourth quintiles. This approach removes all cases of mismeasurement less than \$8,400 ($2 \times 0.1 \times 42,000$). Columns (5) and (6) show that the estimates based on these alternative measures are roughly the same as our baseline estimate.¹⁶

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

4.4. Dynamic difference-in-difference results

The first-difference equation (2) has the advantage of parsimony, but it does not inform whether higher-frequency dynamic effects influence the baseline results. Accordingly, we estimate the dynamic difference-in-difference (DiD) equation

$$New Participant_{i,t} = \mu (Middle_i \times Post_t) + \psi (X_i \times Post_t) + \zeta_i + \rho_t + u_{i,t}, \quad (4)$$

where i and t index investor and week; $Post_t$ indicates if t is greater than the week of the reduction; $New Participant_{i,t}$ indicates if i becomes a robo participant in week t ; ζ_i is an investor fixed effect; and ρ_t is a week fixed effect. Standard errors are clustered by investor and week.

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

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

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

where the coefficients μ_m on the treatment indicator, $Middle_i$, vary non-parametrically by month m , which includes week t . These coefficients represent the difference in outcomes between the middle and upper

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

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

Table 2
Effects of the reduction on middle-class participation.

$Y_i =$	First-difference						Dynamic DiD	
	<i>New Participant_i</i>						<i>New Participant_{i,t}</i>	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Middle_i</i>	15.051*** (1.357)	13.718*** (1.339)	17.987*** (5.116)	10.560** (2.966)	14.516*** (3.060)	15.467*** (1.408)		
<i>Middle_i × Age_i</i>			-0.137 (0.158)					
<i>Middle_i × Post_t</i>							0.823*** (0.067)	0.608*** (0.085)
Measure of Middle	2 nd or 3 rd Quintile	2 nd or 3 rd Quintile	2 nd or 3 rd Quintile	2 nd or 3 rd Quintile	2 nd Quintile	Buffer on Middle	2 nd or 3 rd Quintile	2 nd or 3 rd Quintile
Subsample	Main	Main	Main	Outside Portfolios	Main	Main	Main	Main
Controls	Yes	Yes	Yes	Yes	Yes	Yes		
State FE	No	Yes	Yes	No	Yes	Yes		
Investor FE							Yes	Yes
Week FE							Yes	Yes
Controls × <i>Post_t</i>							No	Yes
State FE × <i>Post_t</i>							No	Yes
R-squared	0.067	0.097	0.097	0.076	0.078	0.098	0.008	0.008
Number of Observations	9,349	9,349	9,349	1,901	7,530	8,982	620,928	620,928

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

$$\text{New Participant}_i = \mu \text{Middle}_i + \psi X_i + \rho + u_i,$$

where *Middle_i* indicates if i belongs to the second (\$1k–\$6k) or third U.S. wealth quintile (\$6k–\$42k), as opposed to the fourth or fifth quintile (>\$42k) that together constitute the reference group; and *New Participant_i* indicates if i becomes a robo participant after the reduction, as opposed to before it. Columns (5)–(6) assess the scope for measurement error by remeasuring *Middle_i*. Column (5) uses an indicator for whether i belongs to the second U.S. wealth quintile, after assigning a missing value to investors from the third U.S. wealth quintile. Column (6) uses an indicator for whether i belongs to the second or third U.S. wealth quintile, after assigning a missing value to investors whose liquid assets are within a 10% buffer of the third quintile. The sample in columns (1)–(3) and (5)–(6) consists all robo participants in the Deposits Dataset. The sample in column (4) consists of all investors in the Portfolio Dataset, which includes robo non-participants and a subset of investors in the Deposits Dataset. State fixed effects cannot be included in column (4) because the state of residence is unobserved for robo non-participants. Columns (7)–(8) estimate the dynamic difference-in-difference (DiD) baseline equation (4),

$$\text{New Participant}_{i,t} = \mu (\text{Middle}_i \times \text{Post}_t) + \zeta_i + \rho_t + \psi_i X_i + u_{i,t},$$

where *Post_t* indicates if t is greater than the week of the reduction; and *New Participant_{i,t}* indexes if i becomes a robo participant in week t , as opposed to the other weeks in our observation window. Column (7) includes week and investor fixed effects. Column (8) additionally interacts *Post_t* with investor controls and a vector of state fixed effects. Wealth consists of liquid assets, defined in Table 1. Wealth quintiles are calculated using the 2016 Survey of Consumer Finances (SCF) dataset. Investor controls X_i are: the log of annual investor income; the investor's age; and an indicator for whether the investor chooses a lower risk tolerance score than that recommended by the robo advisor. Column (4) additionally controls for the log of the outside portfolio's size. Standard errors are clustered by investor (heteroscedasticity robust) in columns (1)–(6), and are two-way clustered by investor and week in columns (7)–(8).

class in each month. We omit the month prior to the reduction, June 2015, such that we interpret the level of μ_m relative to this base period. The parallel trend assumption requires that $\mu_m = 0$ for all months prior to July 2015.

The results in Fig. 4 show that the probability of becoming a robo participant increases sharply and significantly for middle-class investors exactly in the month of the reduction, July 2015. Importantly, the differences between middle and upper-class investors are economically small and statistically indistinguishable from zero before the reduction, providing strong support for the parallel trend assumption. It is thus highly unlikely that pre-trends in middle-class investors' robo participation or other dynamic effects bias the baseline results. Otherwise, such confounding factors would need to coincide exactly with the month of the reduction, which is itself highly unlikely.¹⁷

¹⁷ Specifically, the absence of a pre-trend rules out bias from the launching of new robo products by the two traditional managers described in Section 2.2.


5. Robustness and extensions

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

5.1. Initial deposits

If middle-class investors face minimum-account constraints, then they are more likely to make an initial deposit exactly at the prior

These new products were launched at least two months prior to the reduction (March and May), comfortably before the strong divergence in middle-class investors' behavior. We reiterate that these two traditional managers largely expanded their robo products among existing clients, and so they effectively compete in a different pool than our advisor.



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

$$\text{New Participant}_{i,t} = \sum_{m \neq \text{June 2015}} (\mu_m \times \text{Middle}_i \times \mathbb{1}_{t \in m}) + \zeta_i + \rho_t + \sum_{m \neq \text{June 2015}} (\psi_m \times X_i \times \mathbb{1}_{t \in m}) + u_{i,t}.$$

Subscripts i and t index investor and week. The figure shows the estimated coefficients μ_m , which are interpreted as the average difference in the weekly probability of robo participation between the middle and the upper classes among all weeks in a given month. Brackets are 95% confidence intervals with standard errors clustered by investor and week. The shaded region corresponds to the period after the reduction in account minimum. Time 0 is the month when the reduction was implemented. Time -1 is the reference month (June 2015). The remaining notes are the same as in Table 2.

Fig. 4. Event study. The effects of the reduction in account minimum on the middle class.

minimum of \$5,000, as previously shown in Table 1. Such “bunching” is a hallmark of constrained behavior. By removing these constraints, the reduction should theoretically reduce bunching at the prior minimum.


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

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

5.2. Effects of advertising

We next examine the potential bias in our results due to changes in the advisor's advertising (ad) policy that coincide with the reduction and differentially affect middle-class investors. To address this concern,

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



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

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


Subscripts i and t index investor and week. The figure shows the estimated coefficients μ_m , which are interpreted as the average difference in the first deposit amount between the middle and the upper classes among all weeks in a given month. Brackets are 95% confidence intervals with standard errors clustered by investor and week. The shaded region corresponds to the period after the reduction in account minimum. Time 0 is the month when the reduction was implemented. Time -1 is the reference month (June 2015). We do not include investor fixed effects, since the sample includes only the first deposits. The remaining notes are the same as in Table 2.

Fig. 5. Event study. Bunching behavior by the middle class.

we gather data on the advisor's ad spending during the period around the reduction in account minimum.

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

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



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

Fig. 6. Advertising spending by the robo advisor.

TV advertisements are the dominant method of advertising in the data, accounting for over 99% of spending over the three months before and after the reduction. The small remaining share consists of spending on internet search and display advertisements.¹⁹ Almost all

TV ad spending (98%) is on cable TV advertisements. Considering that addressable advertising on cable TV was less common during that period, it seems unlikely that the advisor specifically targeted individual investors. Nonetheless, there remains the possibility that the advisor broadly targeted investors in specific market areas according to household wealth in those areas.

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

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

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

Lastly, as an additional robustness check, we introduce a different measure of state-level ad spending, denoted $Adv_{w,s,t}$, that accounts for targeting by market areas within states. This measure calculates the average spending in state s in the month of week t , separately for middle-class and upper-class DMAs. Since this measure varies within state-months, we extend our previous specifications from columns (1)–(3) by controlling for $Adv_{w,s,t}$ and $Post_t \times Adv_{w,s,t}$ separately. The

¹⁹ We check whether the dominance of TV advertisements is due to measurement error in the Kantar dataset by tabulating the same statistic using Nielsen's Ad Intel dataset. The Nielsen dataset has a similar form as Kantar's, but it does not contain information on spending across geographic markets and, so, cannot be used for our main analysis. The share of the robo advisor's advertising spending on TV advertisements over the same time period in the Nielsen dataset also equals 99%, suggesting that the main results based on Kantar are not an artifact. Even if the advisor executed unobserved targeted internet advertisements based on wealth, we can still identify the impact of the reduction under the rea-

sonable assumption that internet ad spending targeted at people with a certain wealth level co-varies with television ad spending targeted at regions where average wealth is the same level, which we do observe. Lastly, the marketing literature does not have a clear consensus about the impact of internet advertising on advertiser revenue, as summarized by Gordon et al. (2020), which suggests that unobserved internet advertisements would have minimal impact on our results.

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

Table 3
Advertising as a confounding factor.

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

Note: Standard errors are in parentheses. Point estimates are multiplied by 100. This table estimates a variant of equation (4) that assesses the robustness of the baseline results to a dynamic specification that accounts for changes in advertising that may disproportionately affect middle-class investors. Subscripts i and t index investor and week. The baseline regression equation in columns (1)–(3) is of the form

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

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

results in columns (4)–(6) mirror those obtained from the initial measure.²¹ We therefore conclude that the estimated effect of the reduction does not confound shifts in the advisor's advertising.

5.3. Additional extensions

Gambling. Experimental evidence shows that investors exhibit lower risk aversion in the context of small lotteries (e.g., Bombardini and Trebbi 2012). So, lowering the minimum may induce some account holders to set up an additional, smaller “play account” in which to gamble. Whether this gambling channel drives the results in Table 2 depends on the relative size of a gambling account for the middle versus upper classes. If the optimal size of a gambling account for the upper class lies between \$500 and \$5,000 while the optimal size for the middle class lies below \$500, then Table 2 actually underestimates the effect of relaxing minimum-account constraints. If, on the other hand, the optimal size of a gambling account for the middle class lies between \$500 and \$5,000 while the optimal size for the upper class lies above \$5,000, then the gambling channel may actually drive the main results


in Table 2. We assess this possibility by re-estimating the main specifications from Table 2 after removing the 12% of account holders with multiple accounts from the sample. The resulting estimates in Appendix Table A6 are almost the same as with the full sample, suggesting that gambling motives indeed do not drive the results.

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

Business stealing could affect our results in two key ways. First, the reduction in the account minimum at the advisor may attract *existing* clients from competitor robo advisors. In this case, we would expect to observe a decline in the total number of clients at the competitors after the reduction. The results in Fig. 7 do not support this concern; we actually observe an increase in the number of participants at competitors from the pre-reduction to the post-reduction periods.

Second, the reduction may attract *potential new* clients. In particular, new middle-class robo participants may have planned to invest with a competitor robo advisor during the post-reduction period, but the reduction prompted them to invest with the advisor instead. This concern would imply that the competitors continue growing after the reduc-

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



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

Fig. 7. Participants across robo advisors. Assessing business stealing.

tion, but at a lower rate. The results in Fig. 7 are inconsistent with this concern since they show very little change in growth trends for competitors. We observe an increase in client growth only for the advisor itself, which is consistent with the effects of the reduction.

6. Description of model

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

6.1. Setup

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

6.1.1. Preferences and labor income

Let i index investor. Time is discrete, and t indexes year. For the rest of the exposition, we conserve notation by aligning an investor's age with the year, such that we do not maintain both age and time subscripts. Investors begin their problem at age t_0 . With probability p_t , an investor of age t survives until age $t + 1$, and at age $\bar{T}$ any surviving

investors leave the model. Investors have CRRA preferences over annual flow consumption $C_{i,t}$. Expected lifetime utility is given by:

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

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

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

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

6.1.2. Financial assets

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

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

$$R_{i,t}^P = \beta_i^P F_t + \epsilon_{i,t}^P, \quad (7)$$

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

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

6.1.3. Model solution

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

6.2. Calibration

6.2.1. Portfolio parameters

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

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

$$R_{k,m} = \beta_k F_m + \epsilon_{k,m}, \quad (8)$$

where m indexes month; and $R_{k,m}$ denotes the monthly return on security k , in excess of the risk-free return. We specify F as the three Fama and French (1993) factors with two additional bond factors: the monthly percent changes in the U.S. and the global bond indices from Barclays.²³

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

$$\sigma_{\epsilon,j}^P = \sqrt{w_j^{P'} \hat{\Sigma}_\epsilon w_j^P}, \quad \beta_j^P = w_j^{P'} \hat{\beta}, \quad (9)$$

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

Table 4 compares benchmark characteristics of self-managed and robo portfolios. Robo portfolios feature a 30 pps higher Sharpe ratio for both the middle and upper classes. The higher Sharpe ratio reflects a 2 pps higher expected return and a 11 pps lower idiosyncratic volatility. By construction, the higher expected return stems from greater exposure to priced risk factors, specifically value and bond premia (Appendix Table A13). Bonds, in particular, have a relatively high Sharpe ratio, which undergirds the well-known strategy of risk parity (e.g., Asness et al. 2012). The lower idiosyncratic volatility in robo portfolios reflects how investors in our data mostly hold individual stocks and actively-managed funds. For example, the median self-managed portfolio allocates only 8% to broad-based index-linked ETFs.

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

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

Lastly, in the third step of the calibration, we embed the empirical portfolio parameters in the model. Since age and liquid assets comprise the model's state variables, we project each of the empirical portfolio parameters on the investor's log liquid assets and a fifth-order polynomial in age. Then, we substitute the fitted values into the parameter vector $\{\sigma_{\epsilon,i}^S, \sigma_{\epsilon,i}^A, \beta_i^S, \beta_i^A\}$, using the age and liquid assets of investor i .

6.2.2. Other parameters

Table 5 reports the other parameters. We choose standard preference parameter values of $\gamma = 4$ and $\delta = 0.95$. Appendix Table A19 obtains similar values when structurally estimating these parameters. We follow Cocco et al. (2005) in our calibration of labor income parameters. Accordingly, the deterministic component of income, f_t , is a third-order polynomial in investor age, and the coefficients equal those estimated by Cocco et al. (2005) for their baseline analysis. Similarly, we parameterize $\sigma_v = 0.103$, $\sigma_\epsilon = 0.271$, and $\rho^Y = 0.3\%$. Following Carroll (1997), we parameterize $\phi = 0.12\%$ to match the share of investors with zero total income, based on the 2016 SCF. The remaining parameter values are: $R^f = 0.2\%$, corresponding to the average one-month Treasury yield over 2010–2017; $t_0 = 25$, $\underline{T} = 65$, and $\overline{T} = 100$, all of which are standard; and p_t , which we calculate using the Center for Disease Control's mortality tables (Xu et al. 2020).

6.3. Characterizing investor behavior through policy rules

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

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

Unlike in benchmark models, in our model investors choose between two portfolios. On the one hand, the robo portfolio is preferable because of the advantages just described (e.g., diversification, priced risk exposure). On the other hand, the robo portfolio has an account minimum. Due to this trade-off, investors switch to the robo portfolio after their liquid assets cross some threshold. Panel (d) shows that, under a

Table 4
Summary of self-managed and robo portfolios.

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

Note: This table summarizes portfolios that investors manage themselves (Self-Managed) and portfolios they would receive if they become robo participants (Matched Robo), based on the Portfolio Dataset, which we use to calibrate the model in Section 6. Each observation is a pair of self-managed and robo portfolios. *Sharpe Ratio* is the ratio of expected return to standard deviation of return; *Expected Return* is the expected annual return based on a linear factor model; *Total Volatility* is the standard deviation of return; *Idiosyncratic Volatility* is the standard deviation of the pricing error in the factor model; and *Robo-Sharpe Expected Return* is the counterfactual expected return assuming that the self-managed portfolio invests with the same amount of total volatility as its robo match, which equals the product of the robo Sharpe ratio and the self-managed total volatility. *Return Loss* is the difference between the self-managed portfolio *Expected Return* and its *Robo-Sharpe Expected Return*. The baseline factor model, which is used in this table, is the Fama-French Three Factor Model augmented with two bond factors: the monthly percent changes in Barclays' U.S. and global bond indices (Fama-French with Bond). Expected returns for a portfolio are calculated as the dot product of the portfolio's loading, shown in equation (9), and the risk prices of the factors, reported in Appendix Table A14. Columns (1)-(2) report the mean across portfolios for investors in the second or third U.S. wealth quintiles, and columns (4)-(5) do so for the fourth and fifth quintiles. Columns (3) and (6) test for a difference in mean between matched robo and self-managed portfolios for each wealth class. The sample consists of non-advised portfolios for investors who consult the robo advisor for a free portfolio review. Of these investors, 45% become robo participants. Details on estimating the factor models are in Appendix D. The remaining notes are the same as in Table 1.

\$5,000 account minimum, the threshold ranges between \$25,000 and \$30,000 for most of the investor's working life. This result is broadly consistent with the data, as the median liquid assets of existing middle class robo participants equals \$25,000 (Table 1). Once the investor becomes a robo participant, the risky share jumps (panel (b)) due to the advantages of robo portfolios.

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

7. Effect of the reduction within the model

7.1. Can portfolio differences explain the empirical effect?

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

vestors accumulate wealth much more slowly. For example, nearly 40% of U.S. retail investors over age 40 have less than \$6,000 in liquid assets (Appendix Fig. A4). Our methodology accounts for this empirical regularity.

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

7.1.1. Results

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

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

To gain more confidence in the model, we examine whether it can match other features of the data. Table 6 evaluates the model's fit through three other sets of statistics. Panel (a) reports the pre-reduction and post-reduction share of investors from the middle-class. This share

Table 5
Model parameters.

Parameter	Value	Source
(a) Portfolio Parameters:		
Idiosyncratic Volatility ($\sigma_{e,i}$)	Summarized in Table 4	Portfolio Dataset
Factor Loadings (β_i)	Summarized in Appendix Table A13	Portfolio Dataset
(b) Asset Pricing Factors:		
Market Factor, Mean	0.076	CRSP
Market Factor, Volatility	0.147	CRSP
Fama-French Factors	Appendix Table A14	French
Bond Factors	Appendix Table A14	Barclays
(c) Preferences:		
Coefficient of Relative Risk Aversion (γ)	4	Standard
Discount Factor (δ)	0.95	Standard
(d) Labor Income Parameters:		
Age Profile ($f(x)$)	$0.1682 \cdot x - 0.0323 \cdot x^2/10 + 0.002 \cdot x^3/100$	CGM
Permanent Shock Volatility (σ_v)	0.103	CGM
Temporary Shock Volatility (σ_ε)	0.271	CGM
Correlation with Financial Return (ρ^Y)	0.003	CGM
Probability of Disaster (ϕ)	0.001	SCF
(e) Other Parameters:		
Risk-Free Rate (R^f)	0.002	French
Pre-Retirement Age Range ($[t_0, \underline{T}]$)	[25, 65]	Standard
Range of Survival Rates ($[p(\underline{T}), p(t_0)]$)	[0.865, 0.999]	CDC

Note: This table summarizes the baseline calibration of the portfolio choice model in Section 6. Panel (a) notes the location of the table summarizing portfolio parameters. Panel (b) summarizes asset pricing factors, presenting the mean and volatility of the market factor and notes the location of the tables summarizing the other factors. Factor moments are calibrated using the means and covariances evaluated over the longest available time series over 1960-2017. Panel (c) summarizes preference parameters. Panel (d) summarizes parameters of the labor income process. Note that: a correlation between log labor income and financial returns of $\rho^Y = 0.3\%$ corresponds to a loading of 0.007; the probability of labor income disaster is calculated as the share of investors with zero total income (Carroll 1997), using the 2016 SCF. Panel (e) summarizes other parameters: the risk-free rate corresponds to the average one-month Treasury yield in 2016; investors begin their problem at age t_0 , retire after age $\underline{T}$, and leave the model at $\bar{T} = 100$; and the survival rate corresponds to the probability that an investor of age t survives until age $t + 1$, and it is monotonically decreasing in age. Column (3) reports the source of each value: CRSP denotes the annually-updated stock file from CRSP; French denotes Ken French's website; Barclays denotes the Bloomberg-Barclays aggregate U.S. and unhedged global bond indices; CGM denotes Cocco et al. (2005); CDC denotes the Center for Disease Control's mortality tables; SCF denotes the 2016 Survey of Consumer Finances; and Portfolio Dataset denotes the paper's portfolio dataset summarized in Table 4. Appendix A has details on these data sources.

is essentially a weighted average of the second and third quintiles shown in Fig. 9. Panel (b) reports the risky share for new and existing middle-class participants. The model matches the data well on average, though producing a lower risky share for new middle-class participants. Lastly, panel (c) presents the model-implied stock market participation rate, after adding standard stock market participation costs (Section 7.4.2). Interestingly, the model matches the U.S. stock market participation rate very well.

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

7.2. Welfare gains

We move to the model's second purpose: computing the welfare gains from access to robo advisors. We are particularly interested in the channels that drive any welfare gains and the distribution of these gains. As standard, we measure investor i 's welfare gain by the percent increase in annual consumption under the previous minimum that raises their expected lifetime utility by the same amount as the reduc-


tion. Let $\underline{V}_i$ and $\bar{V}_i$ denote the value functions under the minimums of \$5,000 and \$500, respectively. Appendix C shows that this welfare gain equals:

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

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

7.2.1. Human capital and heterogeneous welfare gains by age

We next examine how welfare gains vary by age. Returning the benchmark comparison group based on actual self-managed portfolios, the results in columns (2)-(4) show that new participants older than 55 gain almost three times as much (1.68%) as new participants aged 35 or younger (0.58%). This finding challenges robo advisors' popular im-



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

Fig. 8. Model mechanics. Policy rules under the previous minimum of \$5,000.

age and their claim to “build our products and services for millennials” (e.g., Hutchins 2020).

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


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

across simulations. We normalize this average by dividing it by the remaining years in the investor's working life.²⁴

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


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

²⁵ This finding does not contradict column (3) of Table 2 because it shows that the magnitude of the welfare gain varies by age, whereas Table 2 shows that the effect on participation does not vary by age. In particular, any investor



Note: This figure plots the share of robo participants from each quintile of the U.S. wealth distribution, based on the portfolio choice model in Section 6. The shares are shown by the open circles. Explicitly, figure shows the distribution across wealth quintiles for participants who find it optimal to participate under the previous minimum (Existing Participants) and for those who find it optimal to participate under the reduced minimum (New Participants). The overall number of existing and new participants is calculated by aggregating investor-level policy functions across the bins of age and liquid assets that define the state space, weighting by share of the 2016 SCF population within each bin. For reference, the bars plot the empirical share of robo participants from each quintile of the U.S. wealth distribution, already shown in Fig. 3. The remaining notes are the same as in Fig. 3.

Fig. 9. Rationalizing the data with a benchmark model of asset allocation.



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

Fig. 10. Understanding why middle-age investors gain more.

Table 6

Rationalizing the data with a benchmark model of asset allocation.

	Model	Data	Source
	(1)	(2)	(3)
(a) Baseline Model, Robo Participation:			
Pre-Reduction Share from Middle-Class	14.9%	12.6%	Deposits Dataset
Post-Reduction Share from Middle-Class	43.1%	28.4%	
(b) Baseline Model, Risky Share:			
Risky Share, Existing Middle-Class	32.3%	31.1%	Deposits Dataset
Risky Share, New Middle-Class	21.4%	30.6%	
(c) Model with Participation Costs:			
Stock Market Participation, Middle-Class	17.3%	19.1%	SCF

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

7.3. Portfolio characteristics that drive the results

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

7.3.1. Inference through counterfactual robo portfolios

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

with a positive welfare gain will choose to participate. So, our model interprets Table 2 as saying that the probability of experiencing a positive gain does not vary by age, even though older investors gain more overall. Consistent with this interpretation, we replicate column (3) of Table 2 in our model and find that the average change in the probability of robo participation for investors under age 50 (+20 pps) is almost the same as for investors above age 50 (+19 pps).

²⁶ An alternative approach would be to remove differences in characteristics one after another. However, this approach might understate or overstate the gains from each characteristic. For example, if we first remove the difference

Table 7
Distribution of welfare gains and channels that drive the gains.

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

Note: This table summarizes the average welfare gain for investors who participate with the robo advisor under the reduced minimum but not under the previous minimum, based on the portfolio choice model in Section 6. Welfare gains are measured by the percent increase in annual consumption under the previous minimum that raises an investor's expected lifetime utility by the same amount as the reduction, as in equation (10). Panel (a) summarizes the average of this statistic for all new middle-class participants in column (1), and columns (2)–(4) do so for various age bins. Panel (b) summarizes the welfare gain when the minimum is reduced on a counterfactual robo portfolio, which helps assess the channels that contribute to the result in (a). The counterfactual portfolios in panel (b.i) constrain the parameters of the robo portfolio for investor i to equal those of the self-managed portfolio for investor i . The constraints are: the same idiosyncratic risk ($\sigma_{e,i}^A = \sigma_{e,i}^S$); the same loading on the market factor ($\beta_{MKT,i}^A = \beta_{MKT,i}^S$); the same loading on value and size factors ($\beta_{HML,i}^A = \beta_{HML,i}^S$, $\beta_{SMB,i}^A = \beta_{SMB,i}^S$); and same the loadings on the U.S. and global bond factors ($\beta_{USB,i}^A = \beta_{USB,i}^S$, $\beta_{GLB,i}^A = \beta_{GLB,i}^S$). The counterfactual portfolios in panel (b.ii) shut down adjustment in the parameters of the robo portfolio ($\beta_i^A, \sigma_{e,i}^A$) by investor characteristics. The adjustments that are shut down are: age, so that robo portfolios only adjust for liquid assets (No Adjustment for Age); liquid assets, so that robo portfolios only adjust for age (No Adjustment for Wealth); and both age and liquid assets, so that robo portfolios do not adjust at all (No Adjustment for Age or Wealth). Panel (c) performs a similar calculation as in panel (a) after specifying a hypothetical self-managed portfolio as the comparison group: Vanguard's Total Stock Market ETF (VTI); and the Vanguard TDF in which an investor of a given age i in 2015 and retirement age 65 would invest. The remaining notes are the same as in Table 2.

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

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

We next consider the welfare effects of the double glide path. We constrain the robo portfolio's asset allocation: to vary by wealth, but not by age; to vary by age, but not by wealth; and to neither vary by age nor by wealth. Panel (b.ii) reports the welfare gain from reducing the minimum on these counterfactual robo portfolios. We find equally smaller gains when removing the adjustment for age (0.34%) or the adjustment for wealth (0.37%). This result suggests that middle-class investors value both adjustments comparably. Column (4) shows that the adjustment by age is particularly important for older investors.²⁷

Our counterfactual analysis suggests that middle-class investors value both basic and sophisticated characteristics of robo portfolios. At a basic level, robo portfolios add value through improving diversification. On its own, this finding may suggest that robo advisors may

²⁷ Personalization by wealth can improve welfare because the two-fund theorem (Tobin 1958) does not hold in our setting. One reason concerns labor income (e.g., Campbell and Viceira 2002). Intuitively, labor income functions like an endowed risky asset. Therefore, the optimal asset allocation, as encoded by the loadings β_i^P , will generically depend on the share of income in total wealth, which, in turn, depends on both the investor's age and wealth. Additionally, the two-fund theorem does not necessarily hold because the account minimum and the implicit borrowing constraint place bounds on the risky share. Lastly, while each household has an optimal robo allocation, the actual robo allocation only approximates the optimal one. Personalization raises welfare insofar as it moves a household closer to her optimal allocation.

offer similar benefits as buying an index fund. However, investors also place significant value on arguably more-sophisticated characteristics, like exposure to additional stock and bond risk factors and a double glide path by age and wealth. Combining these characteristics with improved diversification is important for generating welfare gains, as we shall see from the next approach that directly computes the welfare gain for someone who already invests in an index fund.

7.3.2. Inference through hypothetical comparison groups

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

Panel (c) of Table 7 summarizes the results. A new middle-class participant whose risky portfolio would consist only of VTI without the reduction gains 0.72% in lifetime consumption, or around 90% of the gain for one who would otherwise invest as do the investors in our data. So, while robo portfolios partly resemble a stock index fund, as suggested by Reher and Sun (2019), investors appear to value characteristics of robo portfolios that go beyond this similarity. Otherwise, this hypothetical investor would have experienced no gain.²⁸

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

Based on our findings from this subsection, we conclude that the recent growth in the robo market differs from the concurrent growth observed in index funds and TDFs. While these trends share a connection, the distinction lies in the fact that investors would still gain from access to robo advisors even if they already invested in one of these alternative products. This is because robo portfolios introduce additional layers of

sophistication beyond the features already in TDFs and, especially, in index funds. Hence, even a hypothetical investor with a relatively sophisticated approach of investing in TDFs would nevertheless benefit from accessing robo portfolios. In a broader context, this finding aligns with the advisor's mantra that "everyone deserves sophisticated investment advice", emphasizing that a significant portion of the gains from accessing robo portfolios stems from their more advanced characteristics.

7.4. Additional analysis

7.4.1. Heterogeneity by pre-participation, initial deposit, and wealth

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

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

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

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

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

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

late enough wealth to be willing to invest, relative to new participants from the third quintile. Consequently, they value the reduction in account minimum much more.

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

7.4.2. Model extensions and robustness

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

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

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

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

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

8. Conclusions

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

Second, our results support models of bounded rationality in which retail investors optimally seek professional management given limits on their ability to invest with the same level of sophistication as a professional manager. By gaining access to robo advisors, less-wealthy

investors experience moderate welfare gains. We leave open the question of why households do not invest like a professional, as well as whether these results extend to hybrid wealth management that does not rely fully on automation.

CRedit authorship contribution statement

Michael Reher: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Stanislav Sokolinski:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors have nothing to disclose.

Data availability

The code and data of the article can be found at <https://data.mendeley.com/datasets/3f88g2cb7x/2>.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jfineco.2024.103829>.

References

- Asness, C.S., Frazzini, A., Pedersen, L.H., 2012. Leverage aversion and risk parity. *Financ. Anal. J.*
- Bachas, P., Gertler, P., Higgins, S., Seira, E., 2018. Digital financial services go a long way: transaction costs and financial inclusion. In: *AEA Papers and Proceedings*.
- Bachas, P., Gertler, P., Higgins, S., Seira, E., 2020. How debit cards enable the poor to save more. *J. Finance*.
- Balduzzi, P., Reuter, J., 2019. Heterogeneity in target date funds: strategic risk-taking or risk matching? *Rev. Financ. Stud.*
- Barber, B.M., Huang, X., Odean, T., Schwarz, C., 2022. Attention induced trading and returns: evidence from robinhood users. *J. Finance*.
- Barberis, N., Huang, M., Thaler, R.H., 2006. Individual preferences, monetary gambles, and stock market participation: a case for narrow framing. *Am. Econ. Rev.*
- Bartlett, R., Morse, A., Stanton, R., Wallace, N., 2021. Consumer-lending discrimination in the FinTech era. *J. Financ. Econ.*
- Ben-David, I., Franzoni, F., Kim, B., Moussawi, R., 2022. Competition for attention in the ETF space. *Rev. Financ. Stud.*
- Benzoni, L., Collin-Dufresne, P., Goldstein, R.S., 2007. Portfolio choice over the life-cycle when the stock and labor markets are cointegrated. *J. Finance*.
- Bhattacharya, U., Hackethal, A., Kaesler, S., Loos, B., Meyer, S., 2012. Is unbiased financial advice to retail investors sufficient? Answers from a large field study. *Rev. Financ. Stud.*
- Bianchi, M., 2018. Financial literacy and portfolio dynamics. *J. Finance*.
- Bianchi, M., Brière, M., 2022. Augmenting Investment Decisions with Robo-Advice.
- Bombardini, M., Trebbi, F., 2012. Risk aversion and expected utility theory: an experiment with large and small stakes. *J. Eur. Econ. Assoc.*
- Calvet, L.E., Campbell, J., Sodini, P., 2007. Down or out: assessing the welfare costs of household investment mistakes. *J. Polit. Econ.*
- Calvet, L.E., Campbell, J.Y., Sodini, P., 2009. Fight or flight: portfolio rebalancing by individual investors. *Q. J. Econ.*
- Campbell, J.Y., Cocco, A.F.J., Gomes, F.J., Maenhout, P.J., 2001. Investing Retirement Wealth: A Life-Cycle Model, Risk Aspects of Investment-Based Social Security Reform.
- Campbell, J.Y., Viceira, L.M., 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press.
- Carroll, C.D., 1997. Buffer-stock saving and the life cycle/permanent income hypothesis. *Q. J. Econ.*
- Chalmers, J., Reuter, J., 2020. Is conflicted investment advice better than no advice? *J. Financ. Econ.*
- Chodorow-Reich, G., Nenov, P.T., Simsek, A., 2021. Stock market wealth and the real economy: a local labor market approach. *Am. Econ. Rev.*
- Christelis, D., Jappelli, T., Padula, M., 2010. Cognitive abilities and portfolio choice. *Eur. Econ. Rev.*

- Cocco, J.F., Gomes, F.J., Maenhout, P.J., 2005. Consumption and portfolio choice over the life cycle. *Rev. Financ. Stud.*
- Cole, S., Paulson, A., Shastri, G.K., 2014. Smart money? The effect of education on financial outcomes. *Rev. Financ. Stud.*
- Cookson, A.J., Engelberg, J.E., Mullins, W., 2023. Echo chambers. *Rev. Financ. Stud.*
- Costa, P., Henshaw, J.E., 2022. Quantifying the Investor's View on the Value of Human and Robo-Advice. Vanguard Research Report.
- D'Acunto, F., Prabhala, N., Rossi, A.G., 2019. The Promises and Pitfalls of Robo-Advising, *Review of Financial Studies*.
- D'Acunto, F., Rossi, A.G., 2020. Robo-advising. In: Rau, R., Wardrop, R., Zingales, L. (Eds.), *Palgrave Macmillan Handbook of Technological Finance*.
- Deaton, A., 1991. Savings and liquidity constraints. *Econometrica*.
- D'Hondt, C., De Winne, R., Ghysels, E., Raymond, S., 2020. Artificial intelligence alter egos: who might benefit from robo-investing? *J. Empir. Finance*.
- Egan, M., 2019. Brokers versus retail investors: conflicting interests and dominated products. *J. Finance*.
- Fama, E.F., French, K.R., 1993. Common risk factors in the returns on stocks and bonds. *J. Financ. Econ.*
- Foerster, S., Linnainmaa, J.T., Melzer, B.T., Previtero, A., 2017. Retail financial advice: does one size fit all? *J. Finance*.
- Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., Walther, A., 2021. Predictably unequal? The effects of machine learning on credit markets. *J. Finance*.
- Fuster, A., Plosser, M., Schnabl, P., Vickery, J., 2019. The role of technology in mortgage lending. *Rev. Financ. Stud.*
- Gârleanu, N., Pedersen, L.H., 2018. Efficiently inefficient markets for assets and asset management. *J. Finance*.
- Gennaioli, N., Shleifer, A., Vishny, R., 2015. Money doctors. *J. Finance*.
- Gentzkow, M., 2006. Television and voter turnout. *Q. J. Econ.*
- Gentzkow, M., Shapiro, J.M., 2008. Preschool television viewing and adolescent test scores historical evidence from the Coleman study. *Q. J. Econ.*
- Gomes, F., 2020. Portfolio choice over the life cycle: a survey. *Annu. Rev. Financ. Econ.*
- Gordon, B.R., Jerath, K., Katona, Z., Narayanan, S., Shin, J., Wilbur, K.C., 2020. Inefficiencies in digital advertising markets. *J. Mark.*
- Grinblatt, M., Keloharju, M., Linnainmaa, J., 2011. IQ and stock market participation. *J. Finance*.
- Higgins, S., 2022. Financial Technology Adoption: Network Externalities of Cashless Payments in Mexico.
- Hong, C.Y., Lu, X., Pan, J., 2022. FinTech Adoption and Household Risk-Taking. NBER Working Paper.
- Hong, H., Kubik, J.D., Stein, J.C., 2004. Social interaction and stock market participation. *J. Finance*.
- Hutchins, C., 2020. How Millennials Are Thinking About Their Finances During COVID-19, Technical report, Wealthfront.
- Hvide, H.K., Meling, T.G., Mogstad, M., Vestad, O.L., 2023. Broadband Internet and the stock market investments of individual investors. *J. Finance*.
- Jurek, J.W., Viceira, L.M., 2011. Optimal value and growth tilts in long-horizon portfolios. *Rev. Finance*.
- Kaniel, R., Parham, R., 2017. WSJ category kings - the impact of media attention on consumer and mutual fund investment decisions. *J. Financ. Econ.*
- Linnainmaa, J.T., Melzer, B.T., Previtero, A., 2021. The misguided beliefs of financial advisors. *J. Finance*.
- Loos, B., Previtero, A., Scheurle, S., Hackethal, A., 2020. Robo-advisers and Investor Behavior.
- Lucas, R.E., 1987. *Models of Business Cycles*. Oxford University Press.
- Lynch, A.W., Tan, S., 2011. Labor income dynamics at business-cycle frequencies: implications for portfolio choice. *J. Financ. Econ.*
- Merton, R.C., 1980. On estimating the expected return on the market: an exploratory investigation. *J. Financ. Econ.*
- Moulliet, D., Stolzenbach, J., Bein, A., Wagner, I., 2016. Cost-income ratios and robo-advisory. Deloitte Technical Report.
- Mullainathan, S., Noeth, M., Schoar, A., 2012. The Market for Financial Advice: An Audit Study. NBER Working Paper.
- Nash, A., 2015. It's Time to Kill the Monthly Fee for Small Accounts. Technical report.
- Parker, J.A., Schoar, A., Su, Y., 2021. Retail Financial Innovation and Stock Market Dynamics: the Case of Target Date Funds. NBER Working Paper.
- Philippon, T., 2019. On FinTech and Financial Inclusion. NBER Working Paper.
- Reher, M., Sun, C., 2019. Automated financial management: diversification and clientele effects. *J. Invest. Manag.*
- Rossi, A.G., Utkus, S.P., 2021a. The needs and wants in financial advice: Human versus robo-advising. Working Paper.
- Rossi, A., Utkus, S., 2021b. Who Benefits from Robo-Advising: Evidence from Machine Learning.
- Roussanov, N.L., Ruan, H., Wei, Y., 2020. Marketing mutual funds. *Rev. Financ. Stud.*
- Securities and Exchange Commission (SEC), 2017. Form ADV and IARD Frequently Asked Questions: Item 5.D. SEC Reference Guides.
- Tobin, J., 1958. Liquidity preference as behavior towards risk. *Rev. Econ. Stud.* 25.
- Van Rooij, M., Lusardi, A., Alessie, R., 2011. Financial literacy and stock market participation. *J. Financ. Econ.*
- Vissing-Jørgensen, A., 2002. Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structure. NBER Working Paper.
- Von Gaudecker, H.-M., 2015. How does household portfolio diversification vary with financial literacy and financial advice? *J. Finance*.
- Wealthfront, 2023. Automated Investing Accounts. below \$5,000, Technical report.
- Xu, J., Murphy, S.L., Kochanek, K.D., Arias, E., 2020. Mortality in the United States, 2018. NCHS Data Brief.