Addressing Market Access (Part 1 - Research)
Friday, October 22
12:30 p.m. – 1:15 p.m.

This session presents current research examining the role of financial services in diversity and inclusion in capital markets, including advances in technology, advertising and the availability of financial advice.

Moderator:

Lori Walsh
Vice President, Office of the Chief Economist
FINRA Office of the Chief Economist

Panelists:

Christopher Clifford
Chair, Department of Finance and Quantitative Methods; Phillip Morris Associate of Finance
University of Kentucky

Will Gerken
Real Estate Endowed Associate Professor of Finance
University of Kentucky

Stanislav Sokolinski
Assistant Professor with the Finance and Economics Department
Rutgers Business School

Addressing Market Access (Part 1 - Research) Panelist Bios:

Moderator:

Lori Walsh is Vice President and Deputy Chief Economist and is responsible for management of national strategic and tactical initiatives by providing high quality research, analysis and advice on the economics of securities regulation and policy in support of the Office of the Chief Economist (OCE) and FINRA's mission. The work product informs current rulemaking initiatives, instantiated rules and potential areas of future regulatory actions through analysis of the economic impacts of existing and potential rulemakings. Ms. Walsh has been with FINRA since 2017. In addition to standing in for the Chief Economist in his absence, Ms. Walsh has been leading OCE's efforts to provide strategic and tactical support to FINRA's settlement and litigation activities. Ms. Walsh is also actively involved in several advanced R&D initiatives with Technology and other business units to ensure that FINRA stays a on the cutting edge of technological innovations in the financial markets. Ms. Walsh holds a B.S. in Accounting, an MBA and a PhD in Finance from the Pennsylvania State University. She is also a graduate of FINRA's Leadership Wharton program.

Panelists:

Dr. Chris Clifford is Chair of the Finance and Quantitative Methods department at the University of Kentucky and the Phillip Morris Associate Professor of Finance. His academic research interests focus primarily on non-bank intermediaries such as hedge funds, mutual funds, and investment advisors. His papers have been published at the *Review of Financial Studies*, the *Journal of Financial Economics*, and the *Journal of Finance*. His work has been mentioned in media outlets such as the *Financial Times*, *The Economist*, and *The Wall Street Journal*.

Will Gerken is the Real Estate Endowed Associate Professor of Finance. He has a PhD in Finance from Michigan State University, MS and MBA degrees from Georgia Tech, and BS degrees from West Virginia University. He is a CFA charter holder and serves as the principal contact for the CFA Institute University Affiliation Programs (BS & MSF). Prior to joining the Gatton College of Business, he was an Assistant Professor at Auburn University. His research focuses on financial advisors, financial misconduct, and governance. He has published his research in leading finance journals such as: *Journal of Finance*, *Journal of Financial Economics*, *Journal of Financial and Quantitative Analysis* and *Review of Finance*. He is an associate editor of the *Journal of Corporate Finance*. His research has won best paper awards at leading international conferences. His research has also been featured in the international media such as *The Wall Street Journal*, *Financial Times*, and *NPR* and cited by Securities Exchange Commission.

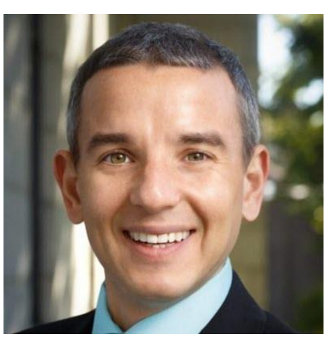

Stanislav Sokolinski is Assistant Professor of Finance and Economics at the Rutgers Business School in Newark and New Brunswick. His current research interests are in financial intermediation and asset management, with a specific emphasis on industrial organization to understand the benefits for investors and market efficiency. He studies the effects of financial advice, the determinants of competition and incentives in the mutual fund industry and explores how new financial technologies affect wealth inequality. Dr. Sokolinski holds a Ph.D. in Economics from Harvard University, a B.A. in Economics and Accounting and M.A. in Financial Economics from Hebrew University. Prior to pursuing his academic career, he advised the Director General of Israeli Ministry of Finance on economic policy and capital market reforms.

Addressing Market Access (Part 1 - Research)

Panelists

- **Moderator**
  - Lori Walsh, Vice President, Office of the Chief Economist, FINRA Office of the Chief Economist
- **Panelists**
  - Christopher Clifford, Chair, Department of Finance and Quantitative Methods; Phillip Morris Associate of Finance, University of Kentucky
  - Will Gerken, Real Estate Endowed Associate Professor of Finance, University of Kentucky
  - Stanislav Sokolinski, Assistant Professor with the Finance and Economics Department, Rutgers Business School

Trust and Race:
Evidence From the Market for Financial Advice

Christopher Clifford, University of Kentucky

Research Question

- Many households do not participate in equity markets
- Financial advisers are overwhelmingly white (and male)
- Trust in financial advisers plays a role in equity market participation

How does racial homophily between adviser and client relate to stock market participation?

Methodology

- FINRA data
  - Track geographic location (branch) of 1.5MM advisors over their career
  - NamePrism Linguistic surname algorithm to determine adviser race
- IRS data
  - Equity market participation at the zip code-income-year level
- Census data (ACS survey)
  - Zip code level demographic and socioeconomic data

Main Takeaways

- **Where a racial match exists e.g.,(at least one black advisor in a predominately black zip code)**
  - Equity market participation rates are 12% higher
    - Control for zip code income and population
    - Effect holds within-zip code changes
  - These effects predominately exist among middle income households ($50k-200k)
    - Relative to their own unconditional participation rates, middle income households are 25% more likely to invest when there is racial homophily
    - Low vs. high income levels?

Implications

- Non-monetary costs can affect participation
- Disparate impact can exacerbate wealth gaps
- Need to consider when hiring and developing programs to establish cross-cultural trust

Where is the Intersection of Madison Avenue and Wall Street?

Advertisement, Local Access to Investment Advice, and Stock Market Participation

Joe Farizo, Will Gerken, and Ge Wu

Research Questions

What are the effects of advertising on stock market participation?

How do these effects interact with disparate levels of income, access to financial advice, and trust?

Methodology

- Data: *Advertising* (Kantar Media), *Advisors,* (SEC Form ADV & FINRA's BrokerCheck), & *Participation* (IRS Statistics of Income)

Main Takeaways

- **Investment advisory firm advertising elasticities are significant**
  - but somewhat smaller in magnitude than those found in other industries (Shapiro, Hitsch, and Tuchman, 2021)

- **Significant income effects**

- **Complementary nature of local access and trust building**

Implications

- Access and awareness of local finance matters

- Under-served neighborhoods may not be able to reap the benefits of promotional campaigns that help encourage broader participation in the markets

Automation and Inequality in Wealth Management

Michael Reher (UCSD) and Stanislav Sokolinski (Rutgers)

Research Question

- Professional wealth managers historically catered to the wealthy.

- Automated wealth managers (i.e. robo advisors) promote the idea of democratizing wealth management
  - Low account minimums and management costs.

- Is this view correct?
  - Do less wealthy participate in robo advice when it becomes more accessible?
  - What are the benefits from access to automated wealth management?

Methodology

Data and Setting

- Novel datasets from a major U.S. robo advisor, Wealthfront.
- Account-level data with demographics and investment activity.
- Experiment: the 2015 reduction in account minimum from $5000 to $500.

Analysis

- Differences in participation with robo advisor across wealth groups.
- Economic model to examine what drives demand for asset management.
- Quantify benefits to investors.

Main Takeaways

- **The reduction mostly affects middle-class households.**
  - Participation with robo advisor increases by 110%.
  - Especially strong effects for lower middle-class.
  - Participation among upper/lower class is unaffected.

- **Investors participate because they struggle to diversify on their own.**
  - Robo-portfolios have Sharpe Ratio of 0.75 against 0.45 for self-managed.
  - Higher expected returns and lower volatility.

- **Welfare gain equates to gain from 4pp higher equity premium.**

Implications

- Robo advice democratizes asset management.
  - Large benefits to investors from improved diversification.

- Why we don't observe even more participation with robo advisors?
  - The robo advice market is still relatively new and it keeps on expanding.
  - Why do some investors keep investing on their own if access to professional wealth management has been improving?
  - Financial education? Lack of information?

Does Automation Democratize Asset Management?*

Michael Reher† and Stanislav Sokolinski‡

February 2021

Abstract

We show that automation affects wealth inequality by giving middle-class households access to asset management. Using novel microdata from a major U.S. automated asset manager (i.e., robo advisor), we study a quasi-experiment in which the advisor suddenly reduces its account minimum by 90%. The reduction relaxes investment constraints on middle-class households and increases the number who participate with the advisor by 110%. Consequently, their expected return on liquid wealth rises by 1-2 pps relative to upper-class households, reflecting a sustained increase in compensated risk. However, automation may not reduce overall wealth inequality, as the reduction does not affect lower-class households.

Keywords: FinTech, Financial Advice, Portfolio Delegation, Inequality

JEL Classification: G11, G24, D3, O3

*This paper was previously circulated under the title "Does FinTech Democratize Investing?". We thank Christoph Aymanns, John Campbell, Francesco D'Acunto, Ben Friedman, Jakub Jurek, Iva Kalcheva, Jun Liu, Matteo Maggiori, Apurva Mehta, Will Mullins, Alessandro Previtero, Enrichetta Ravina, Jonathan Reuter, Alberto Rossi, Andrei Shleifer, Celine Sun, Ross Valkanov, Boris Vallee, anonymous referees, and seminar participants at the Paris FinTech Conference, Toronto FinTech Conference, the New York Fed's FinTech Conference, Georgetown's Finance and Policy Center, UC San Diego, Rutgers Business School, California Corporate Finance Conference, CAFR FinTech Workshop, and Harvard's finance PhD lunch for comments. Ruikai Ji provided excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the position of Wealthfront Inc.

†University of California San Diego, Rady School of Management. Email: mreher@ucsd.edu

‡Rutgers Business School, Newark and New Brunswick. Email: ssokolinski@business.rutgers.edu

1 Introduction

Wealth inequality has soared over the past four decades, in part because the wealthy earn higher financial returns than other households (e.g., Piketty 2014). Wealthy households also have access to a much wider range of investment opportunities, including accounts managed by professional asset managers. Opening such an account typically requires a minimum investment of at least $100,000 (e.g., Pilon 2011), far exceeding the median U.S. household's wealth of $17,000. Against this backdrop, a new class of asset managers that rely on automation (i.e., robo advisors) have promoted the idea of accessible asset management, and they have grown roughly tenfold over the past half-decade.[1]

Whether robo advisors actually improve the financial condition of non-wealthy households is unclear for two reasons. First, it is unclear whether robo advisors even give the non-wealthy access to asset management, versus simply competing with traditional asset managers for the wealthy. Second, even if robo advisors do "democratize" asset management, it is unclear whether this benefits non-wealthy households. For example, participating with a robo advisor may reduce household welfare if, like many asset managers, robo advisors charge high fees, project their own preferences, or face conflicts of interest (e.g., Fama and French 2010; Foerster et al. 2017; Chalmers and Reuter 2020) or if, like other FinTech intermediaries that democratize financial markets, they amplify households' own behavioral biases.[2] On the other hand, automation may enable robo advisors to circumvent these issues and, thus, improve welfare by providing low-fee, diversified, and personalized portfolios (e.g., D'Acunto, Prabhala and Rossi 2019; Loos et al. 2020; Rossi and Utkus 2020).

In this paper, we find that robo advising democratizes professional asset management, and, consequently, it affects inequality in financial returns to the benefit of the middle class (i.e., the moderately wealthy). Our research hypothesis begins with the observation that automation lowers robo advisors' fixed costs relative to traditional asset managers (e.g., Philippon 2019). Therefore, robo advisors can profitably manage smaller portfolios and,

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

[2] The surge in retail trading on the online brokerage Robinhood Markets during the COVID-19 pandemic exemplifies the latter case (e.g., Ben-David et al. 2021; Welch 2020; Barber et al. 2020).

thus, require lower accounts minimums. A lower minimum relaxes investment constraints on middle-class households and, thus, disproportionately increases their participation relative to both the wealthy, who were never constrained, and the poor, who remain constrained. As a consequence of this "asymmetric democratization", financial returns increase for the middle class relative to other households.

Two econometric hurdles make it challenging to test this hypothesis. The first hurdle is limited data: regulatory filings, industry reports, and other readily available sources of data about asset managers do not contain information about the *composition* of households who participate with them, which is central for our research hypothesis. We overcome this challenge by obtaining novel microdata directly from a major U.S. robo advisor. Our dataset includes information on the demographic background, investment activity, and liquid assets, inclusive of retirement accounts, of households who participate with the advisor. This information enables our main investigation into the distributional effects of automation in asset management. Additionally, since our dataset covers a period when the number of robo participants was small, we can examine whether shifts in the composition of robo participants at that time helped spur robo advisors' subsequent growth.

The second econometric hurdle is omitted variables bias. It would be naive to attribute differences in wealth between clients of robo advisors vs. traditional asset managers to lower account minimums, since robo advisors may attract households who vary in other dimensions that correlate with wealth (e.g., technological savviness). We overcome this challenge by studying a quasi-experiment in which the same robo advisor unexpectedly reduces its account minimum from \$5,000 to \$500 in July 2015. This \$4,500 reduction constitutes a large shock for most U.S. households, as it equals 26% of the median U.S. household's liquid assets of \$17,000 at the time. Moreover, it generates a clean source of variation with which to test our research hypothesis because it occurs within the same asset manager, thereby allowing us to hold manager-specific effects fixed.

Our main result is that the reduction democratizes the market for automated asset management by relaxing investment constraints on middle-class households. Graphically, the wealth distribution of robo participants shifts sharply leftward after the reduction, as shown in Figures 1 and 2, while showing no pre-trend in the months leading up to it, as shown in Figure 3. In particular, the share of robo participants from the second and third U.S. wealth quintiles (i.e., "middle class") increases by 107% (16 pps), reflecting a sharp break from trend that is not present among participants from the wealthiest two quintiles (i.e., "upper class"). However, the democratization is asymmetric, in that there is no change in participation among the poorest quintile (i.e., "lower class").

We formalize this graphical intuition through a difference-in-difference analysis that compares the probability of participating with the robo advisor after vs. before the reduction between middle vs. upper-class households. Intuitively, the middle class represents the "treated" group in that it experiences a relaxation of investment constraints due to the reduction. Accordingly, we find that middle-class households are 14 pps more likely to participate with the robo advisor after the reduction, relative to the upper class. Since the middle class was previously underrepresented, this estimate implies that the reduction increases the total number of middle-class robo participants by 110%. These findings do not confound the effect of household demographic characteristics or risk attitude, which we control for in our regressions. Neither do they stem from measurement error in liquid assets, since we obtain the same results from four different measures of middle-class wealth status.

The internal validity of this research design depends on whether the middle class experiences a relaxation of investment constraints. Consistent with this view, the majority of middle-class households who became robo participants prior to the reduction "bunched" their investment right at the previous minimum of $5,000, suggesting that they were constrained by it. After the reduction, however, such bunching immediately disappears, and most new middle-class participants make a previously infeasible investment of less than $5,000. These new middle-class participants were also previously inactive in many financial markets, and, for example, were 43 pps less likely to have participated in the stock market before the reduction, relative to new upper-class participants. This finding again suggests that the previous account minimum imposed a binding constraint on middle-class households' set of investment opportunities.

Based on a wide variety of robustness tests, we find no evidence that the results are driven by channels distinct from a relaxation of constraints, such as targeted advertising, media attention, business stealing from competitors, effects specific to the millennial generation, or gambling motives. For example, we obtain the same results from a panel regression that explicitly allows the middle class to respond differently to the advisor's advertising efforts and media attention, which implies that our research design disentangles the effect of investment constraints from the well-documented impact of visibility on retail investors' behavior (e.g., Kaniel and Parham 2017). Moreover, we find similar results from an aggregated difference-in-difference analysis where observational units are bins of the U.S. population.

Turning to inequality in financial returns, we follow Calvet, Campbell and Sodini (2007) and estimate the reduction's effect on middle-class households' total portfolio return, defined as the expected annual return on liquid assets. We find that the reduction increases the middle class' total return by 1.1 pps relative to the upper class, and this effect is driven by a 13 pps increase in risky share. The increase is strongest among households from the second quintile of the U.S. wealth distribution, as their total return grows by 2.1 pps. These results are robust to a placebo test on existing participants, and we check that they are not driven by various forms of measurement error in risky share or expected return. Neither are our results biased by the possibility that new middle-class robo participants invest more efficiently than other households, since the robo advisor chooses portfolio allocations algorithmically with little room for discretion by participants themselves.

While we cannot make definitive welfare statements without imposing a specific utility function or observing the entirety of households' assets and liabilities, three pieces of evidence suggest that the reduction benefits middle-class households. First, their increase in risky share can be rationalized with plausible coefficients of relative risk aversion (e.g., 7), suggesting that these households do not take excessive risk. Second, the risk they do take is well-compensated, in that 96% of the variance in their robo portfolio returns is spanned by well-known risk factors (e.g., Fama and French 1993). Third, these households hold their positions long enough to realize gains from automated asset management, as only 6% make a withdrawal over our sample period. Moreover, 70% make a subsequent deposit, consistent with the dollar cost averaging approach advocated by practitioners and inconsistent with silent attrition or inertia (e.g., Agnew, Balduzzi and Sunden 2003; Bilias, Georgarakos and Haliassos 2010). Collectively, therefore, we interpret the reduction as a Pareto-improving technological innovation that, by favoring the middle class over both the wealthy and the poor, has an ambiguous effect on overall inequality in financial returns.

In policy terms, governments in both the U.S. and the U.K. have developed retirement plans that aim to increase investing by non-wealthy households (e.g., myRA, OregonSaves, NEST). However, these plans typically offer a limited set of investment options, charge high fees, and provide little personalization.[3] Our results suggest that private asset management can accomplish similar objectives as these programs by using automation to provide portfolios that are both sophisticated and accessible. However, this conclusion comes with two caveats related to external validity. First, aggregating our estimates may overstate the effects of an industry-wide reduction in account minimums, since asset managers compete for the same clients in general equilibrium. That said, we observe little reallocation across robo advisors in our setting, suggesting that such an overstatement may be modest. Second, our quasi-experiment occurs against the backdrop of a rapidly growing market for robo advising (e.g., D'Acunto and Rossi 2020), and so the same research design may yield weaker estimates in, say, an aging economy.

The rest of the paper proceeds as follows. We conclude this section by situating our contribution within the related literature. Section 2 presents an organizing theoretical framework. Section 3 provides institutional background and describes our quasi-experiment. Section 4 describes our data. Section 5 estimates the effect of the reduction on the democratization of the robo market, and Section 6 assesses the robustness of this effect. Section 7 studies the effect on inequality in financial returns. Section 8 studies welfare implications. Section 9 concludes. The online appendix contains additional material.

Related Literature

This paper makes three contributions to the literature. First, we contribute to a nascent literature on robo advisors by showing how automation democratizes professional asset management. Our focus on how robo advisors expand accessibility complements existing research on how they improve diversification, reduce behavioral biases, and increase risky investment

[3] For example, the recently phased-out myRA program only offered investments in U.S. government bonds. OregonSaves charges a management fee of 1%. The U.K.'s NEST pension scheme charges a load of 1.8%, a 0.3% management fee, and can require a minimum contribution.

among investors who already have access to investment professionals (e.g., D'Acunto, Prabhala and Rossi 2019; Loos et al. 2020; D'Hondt et al. 2020; Bianchi and Briére 2020; Rossi and Utkus 2020; Reher and Sun 2019). Like the existing research, our findings contrast with the high fees, underperformance, and misaligned incentives associated with traditional asset managers and financial advisors (e.g., French 2008; Bailey, Kumar and Ng 2011; Christoffersen, Evans and Musto 2013; Del Guercio and Reuter 2014; Linnainmaa, Melzer and Previtero 2021). However, we draw a parallel between robo and traditional advisors, in that both increase stock market participation (e.g., Linnainmaa et al. 2020).

Second, we contribute to a broader literature on new financial technologies (i.e., FinTech) by proposing robo advice as an example of how FinTech affects financial inclusion and wealth inequality. In terms of financial inclusion, this finding complements analogous results in the contexts of app-based payments (e.g., Hong, Lu and Pan 2020), bank deposits (e.g., Bachas et al. 2018; Bachas et al. 2020; Higgins 2020), and mortgage markets (e.g., Fuster et al. 2019; Bartlett et al. 2021; Fuster et al. 2021). In terms of inequality, our empirical results confirm the theoretical prediction of Philippon (2019) that robo advising favors the middle class over both the upper and lower classes. More broadly, this finding exemplifies how FinTech can affect well-documented inequality in financial returns (e.g., Lusardi, Michaud and Mitchell 2017; Campbell, Ramadorai and Ranish 2019; Bach, Calvet and Sodini 2020; Fagereng et al. 2020).

Third, we contribute to a large literature on household finance by empirically characterizing a novel friction that constrains household investment in risky asset markets: account minimums required by asset managers. This friction arises from the supply side and does not directly depend on household characteristics such as preferences (e.g., Barberis, Huang and Thaler 2006), sophistication (e.g., Grinblatt, Keloharju and Linnainmaa 2011; Christelis, Jappelli and Padula 2010), socialization (e.g., Hong, Kubik and Stein 2004), or education (e.g., Cole, Paulson and Shastry 2014; Van Rooij, Lusardi and Alessie 2011).[4]

[4]Based on a calibration, Haliassos and Bertaut (1995) conclude that account minimums may have a quantitatively small effect on household investment, but they do not actually use data on such minimums. Separately, a number of asset pricing models have studied how limited stock market participation may contribute to the equity premium puzzle (e.g., Mankiw and Zeldes 1991, Gomes and Michaelides 2008, or Malloy, Moskowitz and Vissing-Jørgensen 2009). A common parameter in these models is a fixed cost of participation (e.g., Vissing-Jørgensen 2002), and our results show how account minimums can be used to micro-found this parameter. Lastly, the asymmetric democratization that we document provides an

2 Theoretical Framework

We organize our empirical analysis around a framework of delegated risky investment with constraints, which we sketch below. Given our empirical focus, we defer a more formal theoretical treatment to Appendix B.

Setup

1. *Delegated Risky Investment*: Households solve a portfolio optimization problem in which they delegate their risky portfolio to asset managers. This strong preference for delegation is a simplification, but it does match the empirical observation that 70% of stock market participants rely on professional advice, as shown in Appendix Table A3 and corroborated by Guiso and Sodini (2013).[5] The intuition would be the same if, instead, households delegate a fraction ρ of their risky portfolio to asset managers because, say, they do not fully trust them. For the sake of exposition, we discuss the case where ρ approaches one.

2. *Account Minimum*: Asset managers offer a portfolio with risky, net-of-fee return R. However, they can only manage portfolios larger than some account minimum M because they incur a fixed cost of management per portfolio.

3. *Constrained-Optimal Portfolio Choice*: Absent the account minimum, a household would delegate a share $\tilde{\omega}$ of her wealth to the asset manager and receive expected utility $U(\tilde{\omega})$. However, if $\tilde{\omega} < \frac{M}{W}$, then she cannot invest this optimal share, since doing so would result in an investment smaller than the minimum. Instead, she would need to invest the exact share $\frac{M}{W}$ to participate, which she may not do for one of two reasons. First, she may simply lack enough wealth to invest without borrowing (i.e., $W < M$). Second, she may have enough wealth, but participating would require investing such a large share of her wealth that she prefers not to invest at all (i.e.,

alternative perspective to models that predict a positive relationship between FinTech and inequality (e.g., Begenau, Farboodi and Veldkamp 2018; Kacperczyk, Nosal and Stevens 2019; Mihet 2020).

[5] Various theories exist for why households exhibit such a strong preference for delegation (e.g., Gennaioli, Shleifer and Vishny 2015; Gârleanu and Pedersen 2018), but we simply take it as given. According to Guiso and Sodini (2013), only 12% of retail investors make financial decisions without professional assistance.

$U(0) > U\left(\frac{M}{W}\right))$. Thus, the minimum can still constrain households whose wealth exceeds it, as we find empirically.

Effects of Automation

4. *Reduction in Account Minimum*: Technological innovation enables asset managers to automate portfolio management. Automation lowers their fixed per-portfolio costs, and, consequently, they can substantially reduce their account minimum to $M' < M$.

5. *Asymmetric Democratization*: The automation-enabled reduction in account minimum removes constraints on moderately wealthy households, for whom $\frac{M'}{\tilde{\omega}} \leq W < \frac{M}{\tilde{\omega}}$. It does not affect wealthier households, since they were never constrained. Neither does it affect poorer households, since they remain constrained. Thus, the reduction democratizes asset management by enabling middle-class households to participate in it, but this democratization is asymmetric because the poor are unaffected.

6. *Lower Inequality in Returns and Higher Welfare*: Participating in asset management increases moderately wealthy households' risky share, and, therefore, raises their expected financial return relative to the wealthy. By revealed preference, the reduction makes fully-optimizing households better off relative to a counterfactual of holding cash (i.e., $\tilde{\omega} = 0$). In reality, households may not realize these theoretical welfare gains by taking excessive risk, not diversifying their risk, or prematurely liquidating their portfolios. However, our evidence in Section 8 suggests that the automated nature of robo portfolios prevents households from making such mistakes.

7. *General Equilibrium*: In a competitive equilibrium, asset managers innovate to attract moderately wealthy households. Therefore, we interpret the reduction as the result of a productivity "shock" that, in the Romer (1990) sense, is not exogenous in general equilibrium. However, the reduction constitutes an exogenous shock within the portfolio optimization problem described above, and so we can identify its effect on household choice. In particular, the fixed effects in our regressions absorb general equilibrium effects, as we discuss in Section 5.2. Thus, general equilibrium considerations (e.g.,

industry competition) do not affect the results' internal validity, although they may impact external validity in ways outlined our conclusion.

3 Institutional Background

To test the theory outlined above, we study a quasi-experiment in which a major U.S. robo advisor reduces its account minimum. We now describe the U.S. robo advising market, the business model of the particular robo advisor we study, and the reduction implemented by this advisor.

3.1 The U.S. Robo Advising Market

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

At the time of our analysis, Wealthfront managed roughly $3 billion and was the largest standalone robo advisor in the U.S. market, with Betterment and Personal Capital as its nearest competitors. Two traditional asset managers, Vanguard and Charles Schwab, launched robo advising services early in 2015. Both of these services managed more than Wealthfront because they transferred assets from existing, non-robo services. Appendix Table A1 summarizes the largest robo advisors in the U.S. as of July 2015, including their account minimums, assets under management, fees, and provision of traditional, human-based management. Note that Wealthfront stands out as the only robo advisor that relies purely on automation, with no option for a human advisor.

3.2 Robo Portfolio Allocations

Wealthfront, henceforth "the robo advisor", has offered many services throughout its history, including tax loss harvesting, long term financial planning, portfolio lines of credit, and a risk parity fund. Its baseline product, which is most relevant for this paper, is an automatically rebalanced portfolio of 10 ETFs corresponding to 10 asset classes.[6] The portfolio weights are determined by a questionnaire that asks the client several questions about age, liquid assets, income, demographic background, and response to hypothetical investment decisions. The client is then assigned to one of 20 possible risk tolerance scores, which range from 0.5 to 10 in increments of 0.5. Each risk tolerance score uniquely determines a robo portfolio. The portfolio weights solve a problem of optimal asset allocation across the 10 ETFs, taking this score as a parameter. As summarized in Appendix Table A2, portfolios associated with higher risk tolerance scores exhibit higher betas, higher expected returns, and higher proportions of wealth invested in stocks.

Summarizing, the robo portfolios we study conform to most "textbook" recommendations for retail investors (e.g., Malkiel 2015), in that they provide well-diversified risk exposure with more personalization than a generic "60/40" portfolio, but without the complexity often associated with active management. Importantly, robo portfolios are not recommendations, but, rather, they are directly managed by the robo advisor. Consequently, households have little discretion over their portfolio allocations, and so their robo performance will not depend on sophistication (e.g., Grinblatt, Keloharju and Linnainmaa 2011; Christelis, Jappelli and Padula 2010), ability to diversify (e.g., Calvet, Campbell and Sodini 2007), willingness to follow advice (e.g., Bhattacharya et al. 2012), or reluctance to rebalance (e.g., Calvet, Campbell and Sodini 2009).

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

3.3 The 2015 Reduction in Account Minimum

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

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

4 Data

Our core analysis relies on two datasets: a panel dataset covering deposit activity by households who participate with the robo advisor; and a cross-sectional dataset covering all U.S. households. We describe the key features of each dataset here and defer additional

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

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

details to Appendix A. For the rest of the paper, we use the term "robo participant" to describe households who have invested money with the robo advisor, Wealthfront.

4.1 Robo Advising Dataset

The first core dataset contains a weekly time series of deposits with the robo advisor from December 1, 2014 through February 29, 2016. This window straddles the reduction in account minimum, and it marks a formative period in the history of the robo advising market when the number of participating households was still small. We obtained this dataset through a direct query of the robo advisor's internal server, and so we observe the same information as would an analyst working for the advisor. Specifically, we observe the date and size of the deposit, whether the deposit comes from a new participant with the robo advisor, and the following demographic variables about the participating household: annual income; state of residence; householder age; and liquid assets, defined as "cash, savings accounts, certificates of deposit, mutual funds, IRAs, 401ks, and public stocks". The demographic variables are self-reported via the robo advisor's questionnaire and static. Thus, liquid assets may be subject to measurement error from misreporting, but the battery of tests in Section 5.3.1 suggest any such measurement error does not bias our results.

Studying a company-specific dataset has two advantages over publicly available datasets such as the SEC's Form ADV filings, which serve as the basis for many industry reports about the robo market. First, estimates of robo participation growth derived from Form ADV data would be highly imprecise because they can include both inactive clients and "clients" who create a username but never provide the robo advisor with any money.[9] Second, unlike public data, our dataset includes information about a robo participant's wealth. This feature allows us to study investment activity across the wealth distribution, which lies at the heart of our research hypothesis.

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

4.2 Survey of Consumer Finances (SCF)

The second core dataset is the 2016 Survey of Consumer Finances (SCF), which includes financial and demographic information about a representative cross-section of U.S. households. The SCF dataset fulfills two purposes in our setting. First, it allows us to benchmark a household's wealth in our robo advising dataset against the U.S. population and, thus, to estimate the shock's effect on the democratization of the robo market. We respectively use the terms "lower class", "middle class", and "upper class" to describe households from the first, second or third, and fourth or fifth quintiles of the overall U.S. distribution of liquid assets, where liquid assets are calculated to match the definition in our robo advising dataset as closely as possible. The corresponding boundary between the lower vs. middle class is \$1,000 in liquid assets, and the boundary between the middle vs. upper class is \$42,000. The second purpose of the SCF dataset is to impute participation in asset management, the stock market, and homeownership, which we do not observe in our robo advising dataset. Section 6.1.3 describes our imputation methodology at a high level, and Appendix C provides complete details. Appendix Tables A3 and A7-A10 summarize the SCF dataset.

4.3 Summary Statistics

Table 1 compares households who become robo participants after the reduction in account minimum with existing participants. The upper panel shows how these 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. The lower panel restricts the comparison to middle-class households, and it conveys a similar pattern. In particular, the median new middle-class participant's initial deposit of \$2,000 would have been infeasible under the previous account minimum of \$5,000. Indeed, over half of existing middle-class participants invested exactly \$5,000 for their initial deposit, suggesting that they were constrained by the previous account minimum. Interestingly, new middle-class participants are not significantly younger than existing ones, providing suggestive evidence that the reduction works through a relaxation of constraints rather than through, say, technological savviness or other generation-specific effects.

5 Democratization of the Robo Market

We now test our main research hypothesis that the reduction in account minimum increases robo participation by constrained, middle-class households (i.e., "democratization"), which is a necessary first step toward understanding the effect of automation on such households' financial condition. After first providing graphical evidence, we then formalize our identification strategy, report our main results, and assess the magnitude of the effect.

5.1 Graphical Evidence

Three pieces of graphical evidence support the prediction that the reduction democratizes the robo market. First, Figure 1 shows how the wealth distribution of robo participants shifts left after the reduction. This shift reflects how new robo participants are significantly less wealthy than existing ones, as already documented in Table 1.

Second, Figure 2 shows how the leftward shift documented in Figure 1 makes the robo wealth distribution more representative of the overall U.S. wealth distribution (i.e., more "democratic"). Notably, the share of robo participants from the second and third quintiles of the U.S. wealth distribution grows by 107% (16 pps), while the share from the upper two quintiles falls by 18% (16 pps). However, there is a non-monotonic relationship between robo participation growth and wealth, since lower-class households remain nonparticipants. Therefore, consistent with the prediction from Section 2, the reduction asymmetrically democratizes the robo market. Appendix Figure A1a reproduces this pattern using an alternative measure of middle-class status.

Third, Figure 3 shows how the increase in middle-class households' robo participation occurs strikingly and immediately after the reduction, and Appendix Figure A1b confirms its statistical significance. In particular, the sharp jump and absence of a pre-trend in middle-class participation strongly suggests that this increase does not reflect reverse causality. Otherwise, an exogenous shock to middle-class robo participation coinciding exactly with the month of the reduction would have prompted the advisor to reduce its minimum at exactly that time, which seems implausible. More likely, the advisor accurately judged that reducing its minimum would induce such an increase in middle-class participation.

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

5.2 Identification

Begin with the following flexible model of robo participation in period T,

$$Participant_{i,T} = \beta\left(Middle_i \times Post_T\right) + \delta\left(X_i \times Post_T\right) + \mu_i + \tau Post_T + \upsilon_{i,T}, \tag{1}$$

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

Our framework in Section 2 predicts that the reduction affects robo participation among households with moderate levels of wealth because it relaxes constraints on their ability to invest. Equation (1) measures "moderate wealth" using the indicator $Middle_i$. Therefore, under an identification assumption described shortly, the parameter β equals the effect of the reduction on middle-class households' probability of robo participation. Explicitly, β equals the double difference in the probability of becoming a robo participant after vs. before the reduction between middle-class vs. upper-class households.

Two other terms in equation (1) merit discussion. First, μ_i captures time-invariant or otherwise slow-moving characteristics that predispose households to participating with the advisor, such as a certain level of sophistication or trust (e.g., Guiso, Sapienza and Zingales 2008). Since such "affinity" to the advisor increases the probability of participation in any period, we can separately identify the effect of investment constraints because the account minimum changes over time. Second, the interaction between X_i and $Post_T$ captures heterogeneous trends by observed household characteristics. If, for example, younger households are more likely to become robo participants after the reduction for reasons apart from a relaxation of investment constraints, then this effect would be separately captured by the parameter δ. Thus, both of the additional terms in equation (1) capture channels distinct from the investment constraints channel, and they also define how strong the latter must be to induce a change in robo participation. In particular, relaxing constraints increases the probability of robo participation by β, and this effect can push a household over the threshold to eventually participating if the other two channels are strong enough.

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

$$\Delta Participant_i \equiv New\ Participant_i = \beta Middle_i + \delta X_i + \tau + u_i, \tag{2}$$

where $New\ Participant_i$ indicates if household i becomes a robo participant after the reduction; and $u_i \equiv \Delta v_i$. We estimate equation (2) on the set of eventual robo participants, and, therefore, β equals the reduction's effect on the probability of robo participation conditional on eventually participating. This statistic is relevant because, as we formalize in Section 5.4, it directly maps to the mass of participants whom the reduction brings into the robo market and, thus, quantifies the democratization of the robo market.

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

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

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

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

Second, equation (3) could be violated if u_i captures changes in middle-class households' robo participation that are themselves side effects of the reduction, the leading examples of which are media attention and advertising. If middle-class households are more exposed to such media attention or advertising, then the results may confound the effects of these alternative shocks, which work through heightened visibility (e.g., Kaniel and Parham 2017), with the effect of the reduction, which works through a relaxation of investment constraints. We assess the scope for bias from heterogeneous visibility in Section 6.2, and the evidence suggests that such heterogeneity does not bias the estimates.

5.3 Baseline Results

Table 2 reports the results. The estimate in column (1) implies that middle-class households are 22 pps more likely to become robo participants after the reduction in account minimum, relative to upper-class households. After we add household-level control variables and state fixed effects in column (3), the estimated effect equals 14 pps, which we take as our baseline estimate. We reserve Section 5.4 for assessing the magnitude of this effect.

5.3.1 Measurement Error

Our treatment exposure variable, $Middle_i$, may be subject to additive measurement error due to self-reporting. On the one hand, such measurement error introduces attenuation bias, which would tend to bias the estimates toward zero. Similarly, the estimates are biased toward zero if new robo participants overreport their wealth more than existing participants do. On the other hand, measurement error biases the estimates away from zero if new participants underreport their wealth relative to existing participants.[10]

We mitigate this concern by remeasuring $Middle_i$ in three ways, all of which plausibly provide a more precise measure. First, we redefine the middle class exclusively as the second quintile of the U.S. wealth distribution and omit households from the third quintile from the sample. Under this definition, upper-class households would need to underreport liquid assets by at least $36,000 to be misclassified as middle-class. Second, we exclude households whose liquid assets are within a 10% buffer of the boundary between the third and fourth quintiles. This approach removes all cases of mismeasurement that exceed $8,400 ($2 \times 0.1 \times 42,000$). Third, we remove households whose reported wealth class differs from their imputed wealth class, based on the imputation procedure described in Appendix C applied to the SCF dataset. The remaining 61% of households have a well-measured wealth class, in that it accords with what one would predict based on a nationally representative dataset. Note that self-reported wealth class is equally well-measured for middle and upper-class robo participants, as shown in Appendix Table A10, which suggests that the measure of liquid assets in our robo advising dataset is not systematically biased.

The estimates based on these alternative measures of $Middle_i$ all lie between 0.1 and 0.17, as shown in columns (4), (5) and (6) of Table 2. This range straddles our baseline estimate of 0.14, suggesting that it is not biased because of measurement error.

5.4 Magnitude of Effect

We use the estimates in Table 2 to decompose the observed growth rate in the total number of robo participants into the component due to the reduction vs. that due to other forces. The observed growth rate can be directly calculated from the data as

$$g = \frac{New\ Participants}{Existing\ Participants}, \tag{4}$$

[10]Formally, if we mismeasure $Middle_i$ as $\widehat{Middle}_i = Middle_i + \varepsilon_i$, then the estimator for β in a specification of equation (2) without controls is: $\widehat{\beta} = \beta\left(1 - \frac{\mathrm{Var}[\varepsilon_i] + \mathbb{E}[Middle_i \times \epsilon_i]}{\mathrm{Var}\left[\widehat{Middle}_i\right]}\right) + \frac{\mathbb{E}[u_i \times \varepsilon_i]}{\mathrm{Var}\left[\widehat{Middle}_i\right]}$. The term in parentheses captures the effect of attenuation bias. The second term captures bias from differences in misreporting between new and existing participants.

where *New Participants* is the number of households who become robo participants after the reduction; and, analogously, *Existing Participants* is the number who participated beforehand. It will be helpful to rewrite the numerator of equation (4) as

$$New\ Participants = \mathbb{E}\left[New\ Participant_i\right] \times All\ Participants, \tag{5}$$

where $All\ Participants = New\ Participants + Existing\ Participants$ is the sum of new and existing robo participants; and $\mathbb{E}\left[New\ Participant_i\right]$ is the share of all such robo participants who are new. Substituting equations (5) and (2) into equation (4) allows us to express g as

$$g = \frac{\mathbb{E}\left[New\ Participant_i\right]}{1 - \mathbb{E}\left[New\ Participant_i\right]} = \frac{\beta\mathbb{E}\left[Middle_i\right] + \delta\mathbb{E}\left[X_i\right] + \tau}{1 - \left(\beta\mathbb{E}\left[Middle_i\right] + \delta\mathbb{E}\left[X_i\right] + \tau\right)}, \tag{6}$$

which, by definition, is numerically equivalent to the expression in equation (4).

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

$$g^C = \frac{\delta\mathbb{E}\left[X_i\right] + \tau}{1 - \left(\delta\mathbb{E}\left[X_i\right] + \tau\right)} = \frac{\mathbb{E}\left[New\ Participant_i\right] - \beta\mathbb{E}\left[Middle_i\right]}{1 - \left(\mathbb{E}\left[New\ Participant_i\right] - \beta\mathbb{E}\left[Middle_i\right]\right)}, \tag{7}$$

where the two expressions on the right side of equation (7) are equivalent. The first expression highlights how the counterfactual growth rate only depends on factors distinct from the reduction, captured by δ and τ. In particular, the effect of the reduction, β, has been removed, as the second expression makes clear. Our statistic of interest is

$$\eta \equiv g - g^C, \tag{8}$$

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

Table 3 summarizes various calculations of η and of the analogous statistic for growth in middle-class households' robo participation.[11] Interpreting the first row, the baseline

[11]The analogous expression for g^C is: $g^C = \frac{\mathbb{E}[New\ Participant_i|Middle_i=1]-\beta}{1-(\mathbb{E}[New\ Participant_i|Middle_i=1]-\beta)}$.

estimates from Table 2 imply that the reduction increases the overall number of robo participants by 14%, which is driven by a 110% increase in the number of middle-class participants. These values quantify the democratization of the robo market, that is, the mass of all robo participants and middle-class participants brought into the market by the reduction, respectively. In relation to Table 2, the 110% increase in the number of middle-class participants follows from the estimated 14 pps increase in their probability of participation because the middle class was underrepresented before the reduction. The additional estimates in Table 3 imply an increase in the number of middle-class participants between 56% and 129%. While we emphasize distributional rather than aggregate effects in this paper, these results nevertheless suggest that an industry-wide adoption of automation could significantly increase the total number of middle-class households participating in asset management.

6 Robustness

We assess the internal validity of the baseline results in Table 2, which is important given that we will again use our baseline setup to assess the reduction's effect on middle-class households' financial condition in Section 7. Specifically, we: directly assess the investment constraints channel (6.1); evaluate dynamic confounding channels, such as media attention (6.2); evaluate a variety of other specific confounding channels (6.3); and perform a complementary analysis where observational units are aggregates (6.4). The results of all these tests support the baseline results' validity.

6.1 Testing the Constraints Channel

According to the theory from Section 2, the reduction increases middle-class households' robo participation because it relaxes investment constraints imposed by the previous minimum. We test this channel in four ways.

6.1.1 Constrained Investment Behavior: Graphical Evidence

We first graphically inspect whether middle-class robo participants invest in a way consistent with binding investment constraints imposed by the previous account minimum, relative to upper-class participants who constitute our control group. Consistent with this hypothesis, 65% of new middle-class robo participants invest under the previous minimum of $5,000, as shown in panel (a) of Figure 4. Such a small investment would have been infeasible under the previous minimum, and so this behavior suggests that many middle-class households would have preferred to invest under $5,000 prior to the reduction but were constrained.

Indeed, panel (b) shows how 52% of middle-class households who became participants before the reduction invest right at the minimum, a hallmark of constrained behavior. However, this bunching behavior dissipates after the reduction, consistent with a relaxation of constraints. Notably, these patterns are much less pronounced among upper-class households, supporting our difference-in-difference assumption (3) that the change in behavior between the middle vs. upper classes represents the effect of investment constraints.

6.1.2 Constrained Investment Behavior: Regression Evidence

Building on the previous exercise, Table 4 reports the results of regressions that complement Figure 4. The estimate in column (1) implies that new middle-class participants are 30 pps more likely to invest under $5,000 than new upper-class participants, which matches the graphical evidence from Figure 4a. Column (2) implies that middle-class households who became participants prior to the reduction were 25 pps more likely to invest right at the minimum than upper-class participants, but their propensity to do so falls by 32 pps afterward. This finding matches the pre-reduction bunching behavior shown in Figure 4b, which then dissipates after the reduction. These results again support our definition of middle-class households as the "treated" group, in that they experience a relaxation of constraints relative to the upper class.

6.1.3 Financial Inclusion Measures

We next test the investment constraints channel by examining new middle-class robo participants' rate of participation in other financial markets (i.e., financial inclusion). If these households did not own enough assets to overcome the previous account minimum, then they presumably could not have participated in other markets that require a minimum investment or, more generally, a fixed cost. Asset management stands out as a classic example of such a market. Similarly, participating in the stock market without outside assistance involves a fixed cost of acquiring financial knowledge (e.g., Lusardi, Michaud and Mitchell 2017). In real estate, accessing the mortgage market to become a homeowner requires a down payment. If the rate of participation in these other markets is the same for new middle-class robo participants as for new upper-class ones, it suggests that new middle-class robo participants were unconstrained by the previous minimum.

Since we do not observe participation in other financial markets in our robo advising dataset, we use the SCF dataset to impute our three measures of financial inclusion: participation in asset management, in the stock market, and in homeownership. Our imputation methodology is standard, as described in great detail in Appendix C. Summarizing, we predict a given measure of financial inclusion for each household in our robo advising dataset using the subset of variables observed in both the SCF and robo advising datasets and a prediction model. Since we take no prior stance on which prediction model to use, we perform this exercise separately for four different conventional and machine learning models, letting the data inform which model is most appropriate based on out-of-sample performance. Among these models, the best out-of-sample performance comes from a tree-based algorithm called "boosted trees". Therefore, after training and testing the boosted trees algorithm on the SCF dataset, we use it to impute the unobserved variables in our robo advising dataset. We obtain similar results from less-accurate models, like logistic regression.[12]

The results in columns (3)-(5) of Table 4 show that new middle-class robo participants are 18 pps less likely to have participated in asset management, 43 pps less likely to have participated in the stock market more generally, and 15 pps less likely to own their home, relative to their counterparts in the upper class.[13] To alleviate concerns about the use of imputed variables, column (6) of Table 4 measures financial inclusion as the probability of

[12]We test four main predictive models: basic logistic regression, logistic regression with regularization, random forest, and boosted regression trees. For each model, we first train the model and optimally choose the model's hyperparameters using a 10-fold cross-validation procedure. We use 80% of the sample for training and cross-validation purposes, and we choose hyperparameters to maximize the ROC-AUC performance metric. We finally test the performance of the model with the optimal hyperparameter set, using the remaining 20% of the data as a test set.

[13]The difference between estimates in columns (3) and (4) reflects how participation in asset management is uncommon among both the middle and upper classes, as shown in Appendix Tables A7 and A8.

living in a zip code where the share of households receiving dividend income exceeds the median share, and we again estimate a negative coefficient. Together, the results in columns (3)-(6) suggest that new middle-class robo participants demonstrate lower levels of financial inclusion and, thus, are likely to have been constrained by the previous minimum.

6.1.4 Participants from Financially Developed Regions

Lastly, we visually inspect the change in the share of robo participants from each U.S. state to assess whether new robo participants come from less financially developed regions. The result in Figure 5 shows how new robo participants do not live in states associated with strong financial services sectors (e.g., New York), but, rather, in states associated with less financial development (e.g., southern states). This observation further supports the role of the investment constraints channel, albeit suggestively.

6.2 Dynamic Confounding Channels

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

$$New\ Participant_{i,t} = \beta\,(Middle_i \times Post_t) + \alpha_i + \tau_t + u_{i,t}, \tag{9}$$

where i and t index household and week; $Post_t$ indicates if t is greater than the week of the reduction; $New\ Participant_{i,t}$ indicates if i becomes a robo participant in week t, as opposed to the other weeks in our observation window; α_i is a household fixed effect; and τ_t is a month fixed effect. The parameter β now equals the effect of the reduction on middle-class households' probability of robo participation in any given week. This interpretation differs from its counterpart in equation (2), where it equals the cumulative effect over the post-reduction period.

As a first step, we estimate equation (9) as-is and report the results in column (1) of Table 5. The reduction increases the weekly probability of becoming a robo participant by 0.7 pps, or, cumulatively, 22 pps over the 32-week post-reduction period (32×0.007), which is on par with the estimated effect in Table 2. The remainder of this subsection augments equation (9) with additional terms to assess the scope for more specific forms of bias.

6.2.1 Media Attention and Targeted Advertising

We examine whether media attention, advertising, or other changes in visibility bias our results by collecting additional data on news articles from Google News and on blog posts written by the robo advisor itself. Then, we create two variables: $\textit{Monthly News Articles}_t$, defined as the number of news articles about the advisor published in the month of week t, which proxies for media attention; and $\textit{Monthly Advisor Blogs}_t$, defined as the number of blog posts written by the advisor in the month of week t, which proxies for advertising.

Our primary concern is that media attention and advertising around the reduction disproportionately influence middle-class households. We address this concern by interacting $Middle_i$ with the previous two proxies, thus allowing middle-class households to respond differentially to changes in the robo advisor's visibility. The corresponding coefficient of interest in columns (2) and (3) of Table 5 is unchanged. This finding suggests that the baseline results do not confound changes in visibility that disproportionately influence the middle class.

6.2.2 Pre-Trends and Other Dynamic Effects

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

We address this concern by replacing $Post_t$ in equation (9) with a set of indicator variables that equal one if the month of week t is k months before the reduction, denoted $\textit{Months Before}_{t,k}$, or after it, denoted $\textit{Months After}_{t,k}$. The coefficients on the interaction between $Middle_i$ and these indicator variables represent the weekly probability that a middle-class household becomes a robo participant during the indicated month, relative to the reference month of June 2015 (i.e., $\textit{Months Before}_{t,1}$). If the increase in middle-class robo participation coincides exactly with the month of the reduction, then unobserved changes in household demand or industry competition can only bias our results it they, too, coincide exactly with that month, imposing a high hurdle for these alternative explanations.

The results in column (4) of Table 5 show that the probability of becoming a robo participant increases sharply and significantly for middle-class households exactly in the month of the reduction, July 2015, consistent with Figure 3. By contrast, the middle and upper classes remain on parallel trends over the preceding months, as implied by the insignificant coefficients on the interactions with $Months\ Before_{t,k}$. The precise timing of this increase makes it unlikely that pre-trends in middle-class households' robo participation or other dynamic effects bias the baseline results.

6.2.3 Dynamic Response by Millennials

Commentators often portray automated asset management as the investment of choice for millennials, defined as being born between 1981 and 1996. On the one hand, millennials may have not yet accumulated enough wealth to participate with traditional asset managers, and so they instead participate with automated ones. This margin would support the investment constraints channel, although it would imply that the reduction relaxes temporary, life-cycle constraints. However, it is unlikely that the reduction exclusively relaxes such temporary constraints, since new middle-class robo participants would otherwise be significantly younger than existing ones, which is inconsistent with Table 1.

On the other hand, millennials may respond more strongly to any increase in the robo advisor's visibility because of, say, technological savviness. This alternative margin would bias our baseline results. We assess the scope for such bias by interacting $Post_t$ with an indicator for whether household i is a millennial, denoted $Millennial_i$. The highly insignificant coefficient on this interaction term in column (5) of Table 5 implies that millennials do not respond to the reduction for reasons apart from belonging to the middle class. This finding supports the baseline results' validity.

6.3 Other Confounding Channels

6.3.1 Business Stealing

New middle-class robo participants may have planned to invest with a competitor robo advisor during the post-reduction period, but the reduction prompted them to invest with Wealthfront instead. In this case, our results would reflect business stealing rather than democratization of automated asset management. We assess this possibility by using data from the SEC's Form ADV to plot new participants at other standalone robo advisors, namely Betterment and Personal Capital. While the Form ADV data have limitations described in Section 4, they are the best source of data for this exercise, short of having microdata from each major U.S. robo advisor. The results in Figure 6 show very little decline in new participation at Wealthfront's competitors, measured by log change in number of clients, from the pre-reduction to the post-reduction periods. This observation suggests that the reduction indeed expands access to asset management, rather than simply reallocating participants across robo advisors.

6.3.2 Gambling Motives

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

6.4 Aggregated Measure of Democratization

As described in Section 5.4, estimating equation (2) on the set of eventual robo participants quantifies the share of such participants who were brought into the robo market by the reduction and, thus, the extent to which the market itself becomes more democratic. A separate but related question is whether the reduction increases the share of households eligible to become robo participants who indeed participate. In Appendix D, we estimate an aggregated version of equation (9) in which observational units are population segments and, thus, not restricted to eventual participants. The estimates from this aggregated analysis are consistent with our main results, as they imply that the reduction increases the probability that any household eligible to become a robo participant does so.

7 Inequality in Returns

As discussed in our introduction, automated asset management does not necessarily improve middle-class household's financial condition, which we measure using the expected annual return on liquid assets. The answer depends on both the quantity of risk that middle-class robo participants take and their compensation for taking that risk. Accordingly, we first estimate the reduction's effect on risky share and then turn to its effect on returns. We assess the welfare content of these effects in Section 8.

7.1 Measuring Risky Share

Let $\textit{Risky Share}_{i,0}$ denote the unobserved risky share during the pre-reduction period (i.e., $T = 0$). If a household's robo investment is financed by cash-on-hand and we ignore the effects of compounding over our relatively short sample period, then liquid assets before the reduction are approximately equal to liquid assets after the reduction. Therefore, the household's risky share after the reduction equals

$$\textit{Risky Share}_{i,1} = \textit{Risky Share}_{i,0} + \frac{\textit{Robo Investment}_{i,1}}{\textit{Liquid Assets}_{i,1}}, \tag{10}$$

where $\textit{Robo Investment}_{i,1}$ is the value of net deposits by i in the post-reduction period (i.e., $T = 1$). We can then express the change in risky share from the pre-reduction to the post-reduction periods as

$$\Delta \textit{Risky Share}_{i} = \frac{\textit{Robo Investment}_{i,1}}{\textit{Liquid Assets}_{i,1}}. \tag{11}$$

Equation (11) may be subject to two sources of measurement error, but, importantly, such measurement error would tend to bias our results toward zero. First, equation (11) would overstate the increase in risky share if robo investments are financed by liquidating an outside risky position. This scenario is unlikely because liquidation would trigger potentially costly capital gains taxes. Avoiding liquidation through a direct transfer would be infeasible because, by design, the robo advisor can only manage a restricted set of ETFs. More importantly, however, our imputations of stock market participation in Appendix Table A7 imply that any such outside liquidation would be especially unlikely for middle-class robo participants, the bulk of whom were formerly nonparticipants in the stock market. Therefore, the measurement error is negatively correlated with the treatment exposure variable, $Middle_i$, leading to conservative bias.

Second, equation (11) would overstate the increase in risky share if robo investments are financed through an increase in saving. However, as we show in Appendix Table A5, the implied increase in the savings rate is smaller for middle-class households, consistent with the fact that such households spend a larger share of their income on necessities (e.g., Aguiar and Bils 2015). Like in the previous case, the corresponding measurement error in equation (11) is negatively correlated with $Middle_i$, again leading to conservative bias.

7.2 Measuring Total Return

Given the change in risky share in equation (11), we can calculate the change in a household's total portfolio return, defined as the expected annual return on liquid assets. Explicitly,

$$\Delta \textit{Total Return}_i = \Delta \textit{Risky Share}_i \times \textit{Risky Return}_i \tag{12}$$

where $\textit{Risky Return}_i$ is a measure of the expected return on household i's robo portfolio.

Using historical averages to measure expected returns is subject to well-known challenges, and so, following Calvet, Campbell and Sodini (2007), we propose an asset pricing model to estimate the expected return for securities in the robo portfolio. Appendix E describes our methodology in detail. Briefly, for each security k, we estimate

$$\textit{Return}_{k,t} = \beta_k^F F_t + \epsilon_{k,t}^F, \tag{13}$$

where F_t denotes a column vector of pricing factors in month t; β_k^F denotes the respective row vector of loadings; and $Return_{k,t}$ denotes the monthly return on security k in excess of the risk-free return. Our baseline model is the standard CAPM, which we later augment with other well-known models (e.g., Fama-French Three Factor). Given a vector of factor risk prices for model F, a vector of portfolio weights across securities k, and a vector of estimated loadings $\hat{\beta}_k^F$, it is straightforward to compute the expected return on household i's robo portfolio under model F, which we denote by $Risky\ Return_i^F$. This return is net of all fees, including the advisor's management fee, if applicable.

7.3 Effect on Risky Share and Total Return

We estimate the effect on new robo participants' risky share and total return using a similar difference-in-difference setup as in equation (2), the validity of which is strongly supported by the robustness exercises in Section 6. The results in column (1) of Table 6 imply that the reduction increases new middle-class robo participants' risky share by 15 pps. Columns (2) and (3) show how this finding is robust to the inclusion of controls and fixed effects. In column (4), we find evidence of substantial heterogeneity within the middle class. In particular, the increase in risky share is 26 pps for households from the second quintile of the U.S. wealth distribution (i.e., "lower middle class"). Appendix Table A6 verifies the canonical intuition that the increase in risky share is greater for households with a higher subjective risk tolerance.

Turning to the effect on inequality in returns, column (5) shows that middle-class robo participants experience a 1.2 pps increase in total return, relative to the upper class. As with risky share, this average effect masks substantial heterogeneity within the middle class, since total return increases by 2.1 pps for the lower middle class, shown in column (8). This effect is quantitatively large given that many of these households were formerly nonparticipants in the stock market, per Table 4. As a back-of-envelope calculation, a 2 pps increase in total return for a former stock market nonparticipant translates to an 800% higher return, given the risk-free rate of 0.25 pps in 2016 ($\frac{2.25-0.25}{0.25}$). Compounding this increase over time implies a potentially large impact on wealth accumulation, given the persistent investment behavior of new middle-class participants that we document in Section 8.3.

7.3.1 Robustness of Effect on Risky Share and Total Return

Table 7 reports the results of a variety robustness tests. Columns (1)-(3) show that the baseline results are robust to the choice of asset pricing model F used to estimate expected return and to using realized return over a three-year period, as in column (4). Column (5) performs a placebo test on the set of households who participated with the advisor before the reduction, which yields insignificant estimates and, thus, supports the baseline results. Column (6) restricts the sample to former nonparticipants in the stock market. The estimates are similar to the baseline ones, supporting the previous claim that the reduction has a large relative effect on such households' total return. Columns (7) and (8) respectively report results based on the subsamples of households with no deposit outflows and with nontaxable, retirement accounts (e.g., IRAs). We interpret both sets of households as long-term investors: in the first case, they do not withdraw funds; and, in the second case, they face high liquidation costs from early withdrawal penalties. The similar estimates suggest that the baseline effect persists over a long horizon, which we corroborate in Section 8.3.

8 Welfare Implications

Before concluding, we examine how the reduction affects households' welfare, relative to a counterfactual of holding their robo investment in cash. To make progress on this question, we evaluate three aspects of new middle-class robo participants' portfolios: the quantity of risk taken; the share of that risk that is compensated; and the length of time over which they can realize the gains from taking compensated risk.

8.1 Excessive Risk

First, we examine whether new robo participants take on an "excessive" quantity of risk by benchmarking their risky share against the optimal share recommended by the classic Merton (1969) formula. This Merton-optimal risky share equals, using our notation,

$$\textit{Risky Share}_i^* = \frac{1}{\gamma}\frac{\textit{Total Return}_i}{\textit{Variance of Return}_i}, \tag{14}$$

where, as in Section 7.2, $Total\ Return_i$ is the expected return on the household's robo portfolio, net of fees and the risk-free rate; $Variance\ of\ Return_i$ is the variance of this return; and γ is the household's coefficient of relative risk aversion. We then define a household's excessive risky share as

$$Excessive\ Risky\ Share_i = Risky\ Share_i - Risky\ Share_i^*, \tag{15}$$

where $Risky\ Share_i$ is defined as in equation (10), after omitting time subscripts because all variables correspond to the post-reduction period.[14]

The upper panel of Table 8 summarizes new robo participants' excessive risky share. Focusing on middle-class participants in column (1), the median household's risky share falls 10 pps short of the level recommended by equation (14) when parameterizing $\gamma = 5$, as in many life cycle models (e.g., Campbell et al. 2001). In fact, the median risky share matches the Merton-optimal share under a higher coefficient of relative risk aversion of $\gamma = 7$, which lies toward the upper end of reasonable parameter values (Mehra and Prescott 1985). Therefore, new middle-class robo participants do not appear to be taking excessive risk. If anything, they appear to underinvest relative to the recommendations of benchmark models. An even stronger statement holds for the small minority of upper-class robo participants who enter the stock market, based on the results in column (2).

8.2 Underdiversification

Second, the Merton (1969) model assumes households diversify away uncompensated risk, and so the results from the previous subsection may be misleading if new middle-class robo participants are not compensated for the risk they take (i.e., underdiversified). We assess this possibility by calculating the share of variance in a household's robo portfolio that is compensated according to asset pricing model F, denoted $Compensated\ Risk\ Share_i^F$ and defined explicitly in Appendix E.

The middle panel of Table 8 summarizes the results. For new middle-class robo partic-

[14] We can only calculate equation (15) for new robo participants who are imputed to have not participated in the stock market beforehand since, for such households, the change in risky share from equation (11) also equals the level of risky share from equation (10). Appendix Table A7 summarizes stock market participation.

ipants, the share of risk that is compensated ranges from 0.64 based on the CAPM to above 0.95 when expanding the set of factors to include the three Fama-French factors, and, in the bottom row, two additional U.S. and global bond factors. This observation implies that robo portfolios do not contain substantial idiosyncratic risk but, rather, expose households to priced risk factors.[15] Notably, behavioral biases do not inhibit diversification because robo portfolio allocations are determined solely by algorithm.

8.3 Persistence of Investment

Third, we evaluate whether new robo participants hold their positions long enough to realize the gains from their robo investment. We are particularly concerned that our findings are driven by a "novelty" effect, where households initially experiment with a new financial service but eventually choose to abandon it. We address this concern by calculating the probability that new middle-class robo participants close their account, withdraw funds, or, on the contrary, make a subsequent deposit.

The results in the bottom panel of Table 8 suggest that new middle-class robo participants exhibit remarkably persistent investment behavior. For example, 97% do not close their account over our sample period. This rate mirrors the 98% non-closure rate among new upper-class participants. Moreover, 94% of new middle-class robo participants make no subsequent withdrawals, and 98% make no withdrawal larger than 20% of their initial deposit. We can also measure persistence by subsequent deposit behavior. Accordingly, 70% of new middle-class robo participants make a subsequent deposit. Such subsequent deposit-making resembles the "dollar cost averaging" strategy commonly advocated by practitioners, which Brennan, Li and Torous (2005) show is optimal for risk averse investors.

[15]We cannot directly rule out correlation between a household's robo portfolio return and nonfinancial income. However, the broad exposure provided by the ETFs in robo portfolios shown in Appendix Table A2 makes it unlikely that households are overexposed to, say, specific labor markets or local business cycles (e.g., Ivković and Weisbenner 2005). Moreover, since the advisor assigns older households a lower risk tolerance score, Appendix Table A2 implies that the allocation to equities decreases in age, consistent with portfolio choice models featuring labor income risk (e.g., Viceira 2002).

8.4 Discussion

The results in Table 8 suggest that the reduction improves welfare for new middle-class robo participants through a non-excessive, sustained increase in compensated risk and, thus, total return. This conclusion is surprising in light of the well-documented underperformance and high fees associated with traditional asset managers and investment mistakes made by retail investors, as referenced in the introduction. However, not all forms of FinTech that democratize financial markets necessarily improve welfare. For example, access to discount brokerage accounts or fractional shares may harm households if they subsequently make poor investment decisions, per footnote 2. In another example, automated advice that is decoupled from asset management may not improve long-term welfare if households do not heed it (e.g., Bhattacharya et al. 2012). Our setting is unique relative to these two examples because the advisor uses automation to both construct and manage a household's portfolio.

The analysis in this section also comes with the caveat that we have specified welfare relative to the particular counterfactual of holding one's robo investment in cash. Alternatively, households may have financed their robo investment by reducing consumption or by borrowing. In the former case, we would tend to understate the welfare implications of the reduction insofar as households do not over-save (e.g., Lusardi and Mitchell 2011; Scholz, Seshadri and Khitatrakun 2006). The latter case is theoretically possible but empirically unlikely given recent evidence that households do not increase borrowing in response to expanded savings opportunities (e.g., Medina and Pagel 2020). Lastly, our organizing theory from Section 2 begins with the idea that households strongly prefer to delegate their portfolio, and we do not account for any nonpecuniary gains from delegation (e.g., peace-of-mind).

9 Conclusion

We found that automation affects inequality in financial returns by bringing middle-class households into the market for asset management. To do so, we studied a large and unexpected reduction in the required account minimum by a major U.S. automated asset manager (i.e., robo advisor). By relaxing constraints on their ability to invest with a professional asset manager, the reduction increases the number of robo participants from the middle segments of the U.S. wealth distribution by 110%. Consequently, these households' expected return on liquid wealth increases by 1-2 pps relative to wealthier households. This increase in return does not reflect excessive or uncompensated risk-taking, nor does it appear to be short-lived. However, the reduction has no impact on the poorest fifth of U.S. households. Therefore, we interpret the reduction as a Pareto-improving technological innovation with ambiguous effects on overall wealth inequality.

From the policy perspective, a number of governments have developed retirement plans for non-wealthy households, with mixed rates of success (e.g., myRA, OregonSaves, NEST). Our results suggest that private asset managers can themselves provide the non-wealthy with retirement plans using automation, without the need for government intervention. However, this conclusion comes with two caveats related to external validity. First, we hold general equilibrium effects fixed in our research design. Accounting for such effects may attenuate the impact of a competitive, industry-wide reduction in account minimums. Second, we study a period of rapid growth in the robo market, and the impact of automated asset management may be weaker in, say, aging or developing economies. Future research may make progress on these questions of external validity by studying automated asset management in a quantitative model. Such a model may also inform how robo advising might be designed to benefit the poorest.

References

Agnew, J., Balduzzi, P. and Sunden, A.: 2003, Portfolio Choice and Trading in a Large 401(k) Plan, *American Economic Review* .

Aguiar, M. and Bils, M.: 2015, Has Consumption Inequality Mirrored Income Inequality?, *American Economic Review* .

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

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

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

Bailey, W., Kumar, A. and Ng, D.: 2011, Behavioral Biases of Mutual Fund Investors, *Journal of Financial Economics* .

Barber, B. M., Huang, X., Odean, T. and Schwarz, C.: 2020, Attention Induced Trading and Returns: Evidence from Robinhood Users.

Barberis, N., Huang, M. and Thaler, R. H.: 2006, Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing, *American Economic Review* .

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

Begenau, J., Farboodi, M. and Veldkamp, L.: 2018, Big Data in Finance and the Growth of Large Firms, *Journal of Monetary Economics* .

Ben-David, I., Franzoni, F., Kim, B. and Moussawi, R.: 2021, Competition for Attention in the ETF Space, *NBER Working Paper* .

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

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

Bilias, Y., Georgarakos, D. and Haliassos, M.: 2010, Portfolio Inertia and Stock Market Fluctuations, *Journal of Money, Credit, and Banking* **42**.

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

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

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

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

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

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

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

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

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

Christoffersen, S., Evans, R. and Musto, D.: 2013, What Do Consumers Fund Flows Maximize? Evidence from Their Brokers Incentives, *The Journal of Finance* .

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

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

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

Del Guercio, D. and Reuter, J.: 2014, Mutual Fund Performance and the Incentive to Generate Alpha, *Journal of Finance* .

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

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

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

Fama, E. F. and French, K. R.: 2010, Luck versus Skill in the Cross-Section of Mutual Fund Returns, *Journal of Finance* .

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

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

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

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

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

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

Gomes, F. and Michaelides, A.: 2008, Asset Pricing with Limited Risk Sharing and Heterogeneous Agents, *Review of Financial Studies* .

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

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

Guiso, L. and Sodini, P.: 2013, Household Finance: An Emerging Field, *Handbook of the Economics of Finance*, Vol. 2, Elsevier.

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

Higgins, S.: 2020, Financial Technology Adoption.

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

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

Ivković, Z. and Weisbenner, S.: 2005, Local Does as Local Is: Information Content of the Geography of Individual Investors' Common Stock Investments, *The Journal of Finance* .

James, G., Witten, D., Hastie, T. and Tibshirani, R.: 2013, *An Introduction to Statistical Learning*, Springer.

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

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

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

Linnainmaa, J. T., Melzer, B. T., Previtero, A. and Foerster, S.: 2020, Investor Protections and Stock Market Participation: An Evaluation of Financial Advisor Oversight.

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

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

Lusardi, A. and Mitchell, O. S.: 2011, Financial Literacy and Planning: Implications for Retirement Well-being, *in* A. Lusardi and O. S. Mitchell (eds), *Financial Literacy: Implications for Retirement Security and the Financial Marketplace.*

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

Malloy, C. J., Moskowitz, T. J. and Vissing-Jørgensen, A.: 2009, Long-Run Stockholder Consumption Risk and Asset Returns, *The Journal of Finance* .

Mankiw, G. N. and Zeldes, S. P.: 1991, The Consumption of Stockholders and Nonstockholders, *Journal of Financial Economics* .

Medina, P. and Pagel, M.: 2020, Does Saving Cause Borrowing?

Mehra, R. and Prescott, E. C.: 1985, The Equity Premium: A Puzzle, *Journal of Monetary Economics* .

Merton, R. C.: 1969, Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case, *The Review of Economics and Statistics* .

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

Murphy, K. P.: 2012, *Machine Learning: A Probabilistic Perspective*, MIT press.

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

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

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

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

Romer, P. M.: 1990, Endogenous Technological Change, *The Journal of Political Economy* .

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

Scholz, J. K., Seshadri, A. and Khitatrakun, S.: 2006, Are Americans Saving "Optimally" for Retirement?, *Journal of Political Economy* .

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

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

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

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

Viceira, L. M.: 2002, Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income, *Journal of Finance* .

Vissing-Jørgensen, A.: 2002, Limited Asset Market Participation and the Elasticity of Intertemporal Substitution, *Journal of Political Economy* .

Welch, I.: 2020, The Wisdom of the Robinhood Crowd, *NBER Working Paper* .

Figures and Tables

Figure 1: Shift in Wealth Distribution of Robo Participants

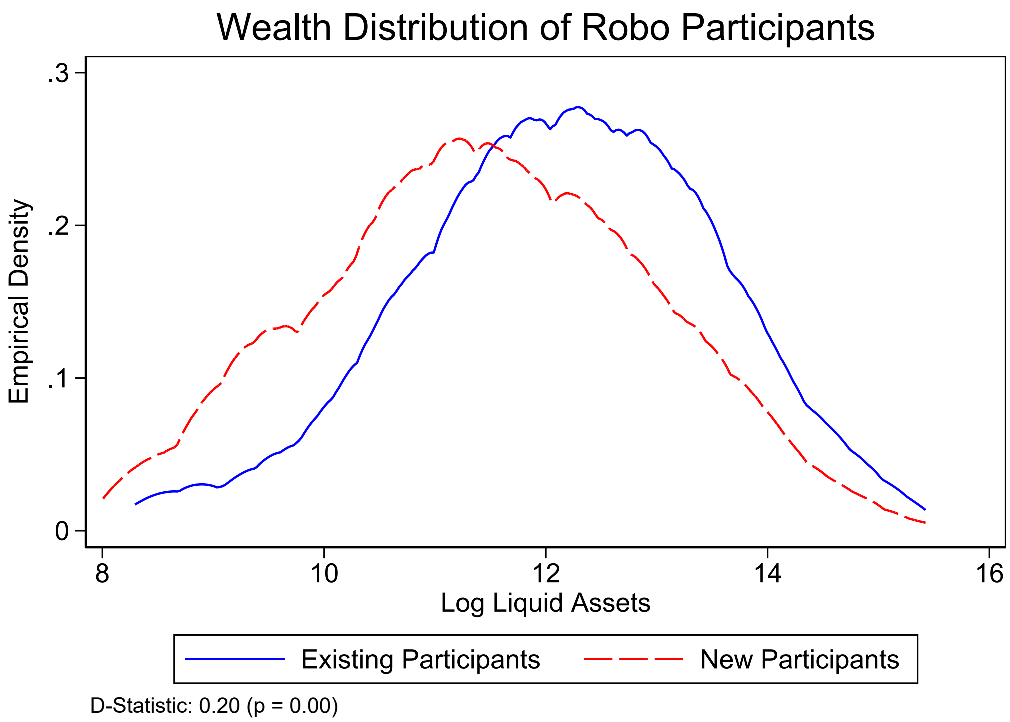

Note: This figure plots the distribution of log liquid assets among households who participated with the robo advisor before the reduction in account minimum (Existing Participants) and who become robo participants after the reduction (New Participants). Liquid assets are defined in Table 1. The distribution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

Figure 2: Change in Representativeness of Robo Wealth Distribution

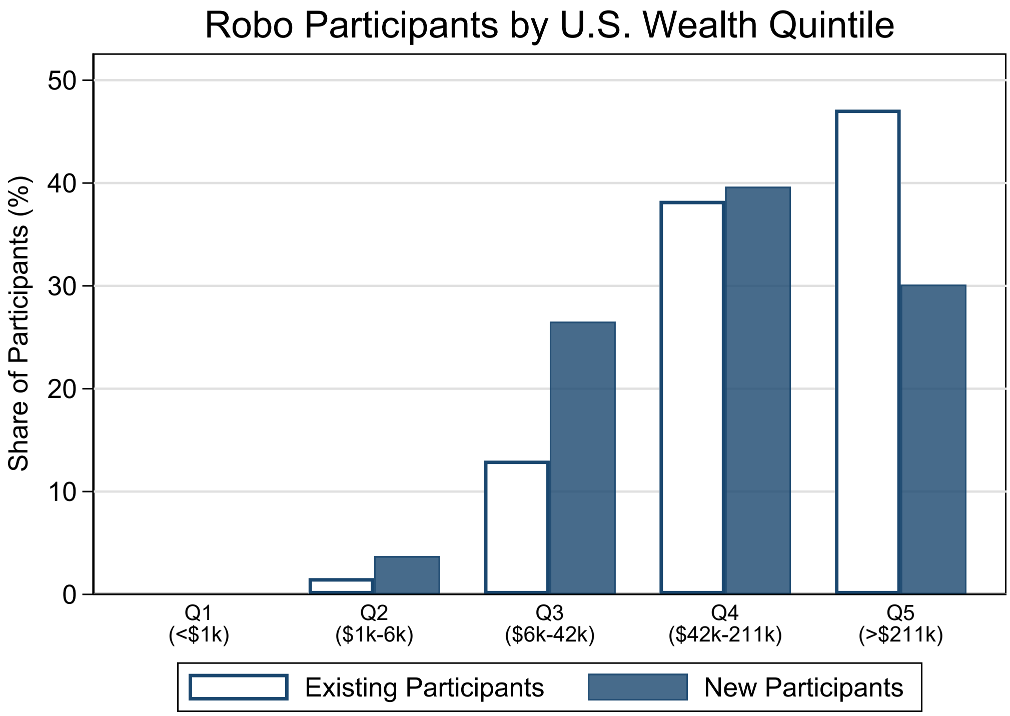

Note: This figure plots the share of robo participants from each quintile of the U.S. wealth distribution. The share is calculated separately for households who participated before the reduction in account minimum (Existing Participants) and who become participants after the reduction (New Participants).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

Figure 3: Pre-Trends in Robo Participation by the Middle Class

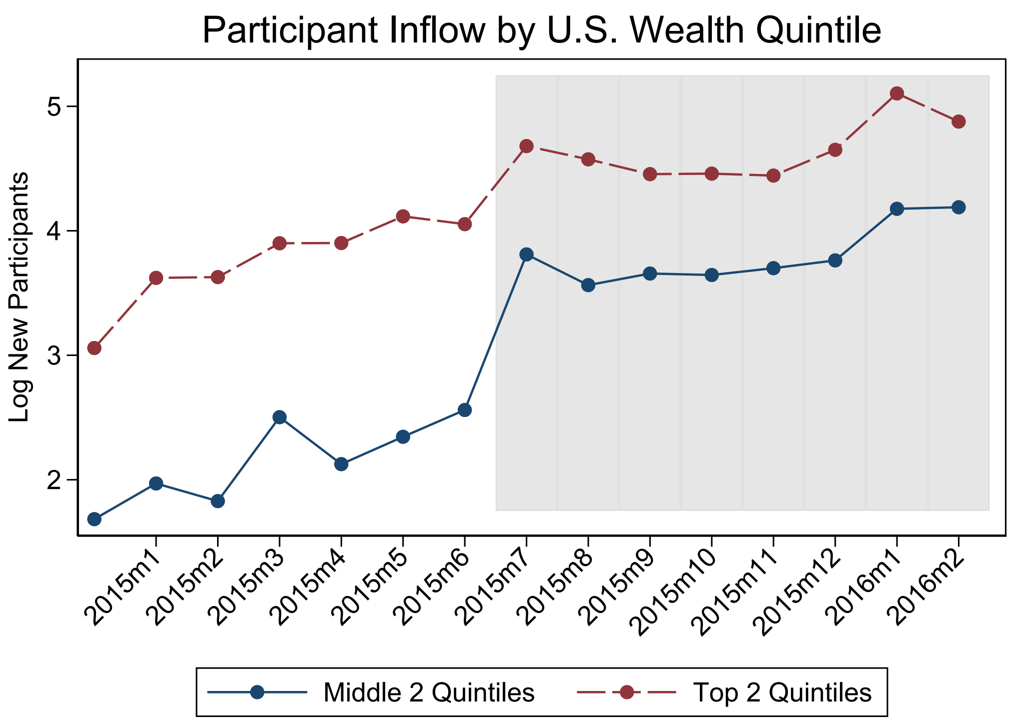

Note: This figure plots the log number of new robo participants from the second and third quintiles (Middle 2) and fourth and fifth quintiles (Top 2) of the U.S. wealth distribution, averaged across weeks in each month. Wealth consists of liquid assets, defined in Table 1. Wealth quintiles are calculated using the Survey of Consumer Finances (SCF) dataset. The shaded region corresponds to the period after the reduction in account minimum.

Figure 4: Constrained Investment Behavior by the Middle Class

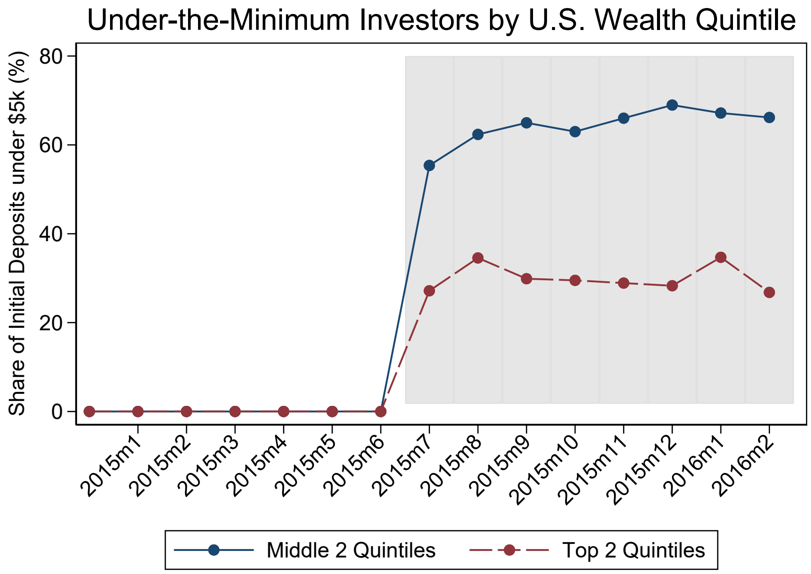

(a) Initial Deposits under $5,000

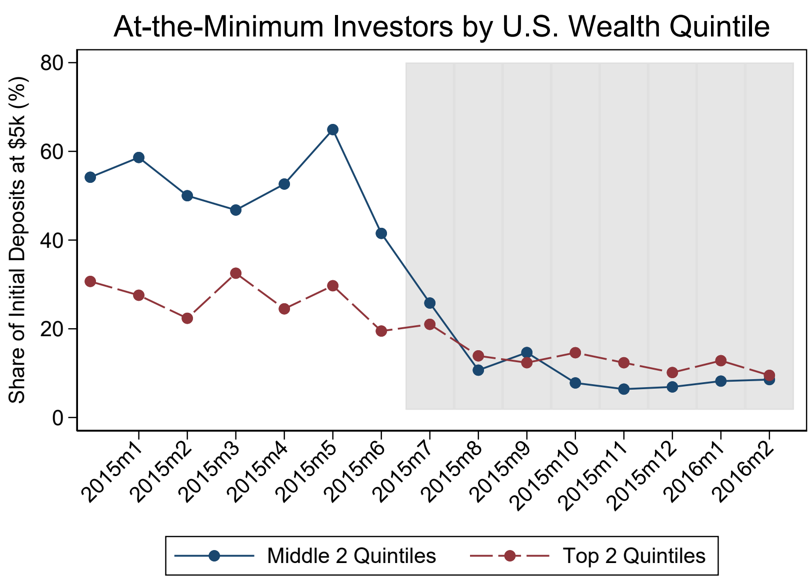

(b) Initial Deposits at $5,000

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

Figure 5: Change in the Geographic Distribution of Robo Participants

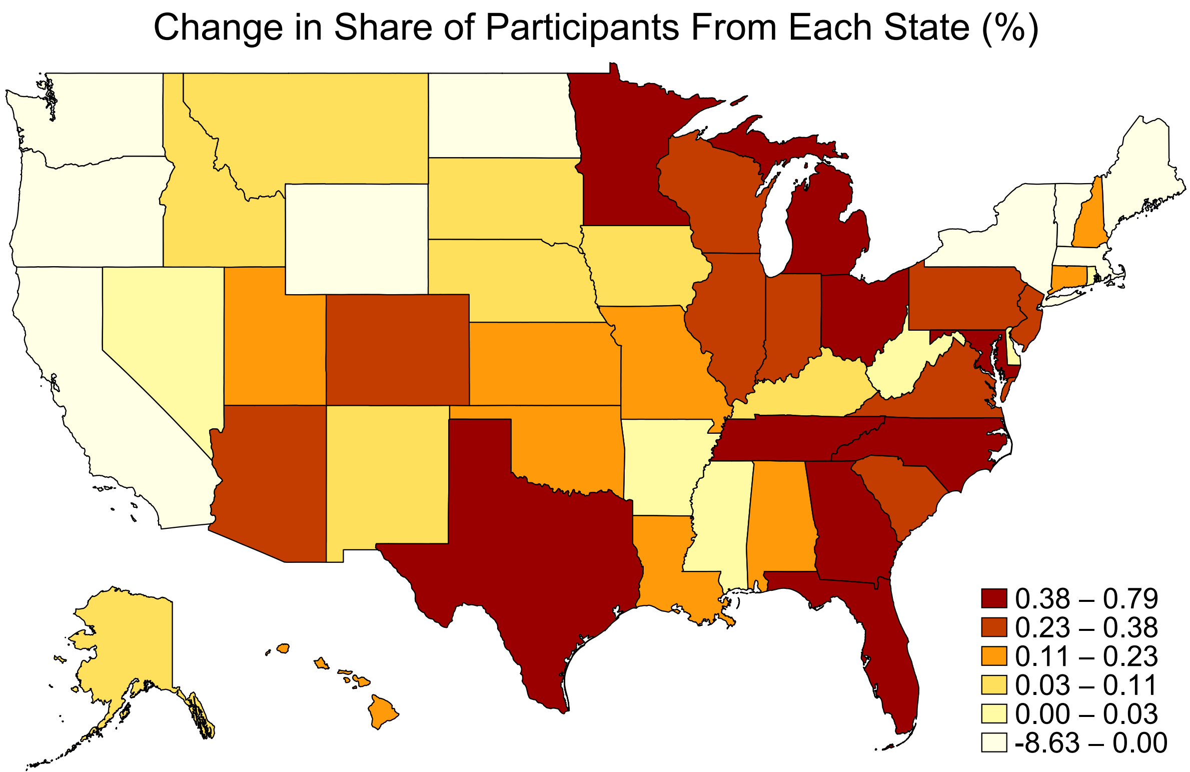

Note: This figure plots the change in the share of robo participants from each U.S. state. The change is from the pre-reduction period (December 1, 2014 to July 7, 2015) to the post-reduction period (July 7, 2015 to February 29, 2016).

Figure 6: Reallocation of Participants across Robo Advisors

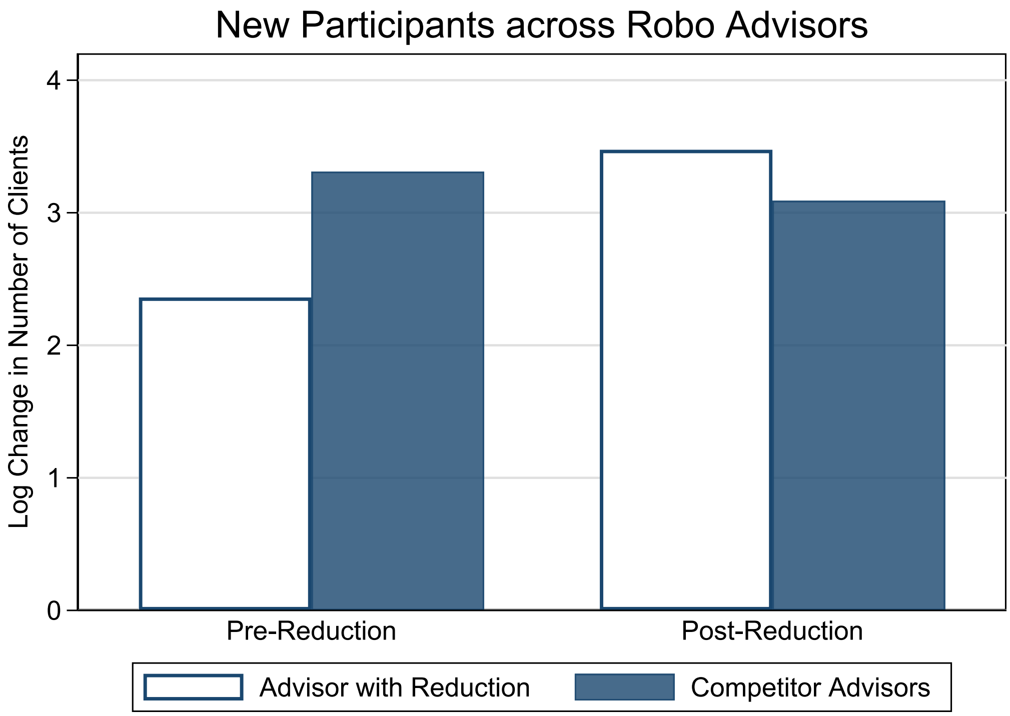

Note: This figure plots the log change in the number of clients across robo advisors, in thousands, which assesses whether the reduction increases robo participation or simply reallocates robo participants across advisors. The change is calculated separately for the robo advisor that reduced its account minimum, Wealthfront, and for its competitors combined. The left two columns plot this change over the pre-reduction period (Q4, 2014 to Q2, 2015), and the right two columns plot this change over the post-reduction period (Q2, 2015 to Q1, 2016). Data are from the SEC's Form ADV. Competitors are defined as Betterment and Personal Capital, since Schwab's and Vanguard's robo advising services do not file a separate Form ADV. The SEC defines clients to include investors who have not compensated their advisor. Advisors do not file a Form ADV every quarter, and so we use the nearest available observation when the advisor does not file a form ADV in a quarter.

Table 1: Summary of Robo Participants

	Existing Participants			New Participants			
	Mean	Standard Deviation	Median	Mean	Standard Deviation	Median	Difference in Mean
All Households:							
Liquid Assets$_i$ (‘000)	436.44	660.82	200	265.21	480.25	100	-171.22
							(0.000)
Income$_i$ (‘000)	157.36	110.67	130	116.17	95.9	90	-41.18
							(0.000)
Initial Deposit$_i$ (‘000)	33.68	94.54	10	22.56	72.61	5	-11.12
							(0.041)
Age$_i$	35.79	8.72	34	35.4	9.97	33	-0.39
							(0.000)
High Risk Tolerance$_i$	0.15	0.35	0	0.14	0.34	0	-0.01
							(0.146)
Middle$_i$	0.15	0.35	0	0.3	0.46	0	0.156
							(0.000)
Middle Class:							
Liquid Assets$_i$ (‘000)	23.23	11.68	25	19.71	11.36	18	-3.527
							(0.000)
Income$_i$ (‘000)	92.86	62.21	80	67.14	42.52	60	-25.720
							(0.000)
Initial Deposit$_i$ (‘000)	7.6	5.34	5	4.95	12.58	2	-2.652
							(0.000)
Age$_i$	30.33	6.33	29	30.04	7.07	28	-0.293
							(0.339)
High Risk Tolerance$_i$	0.13	0.34	0	0.12	0.32	0	-0.012
							(0.446)
Number of Existing Participants: 4,366							
Number of New Participants: 5,336							

Note: P-values are in parentheses. This table summarizes households who participated with the robo advisor before the reduction in account minimum (Existing Participants) and who become participants after the reduction (New Participants). Subscript i indexes household. *Liquid Assets*$_i$ is the sum of cash, savings accounts, certificates of deposit, mutual funds, IRAs, 401ks, and public stocks, in thousands of dollars. *Income*$_i$ is annual household income, in thousands of dollars. *Initial Deposit*$_i$ is the value of the household's initial deposit, in thousands of dollars. *Age*$_i$ is the householder's age. *High Risk Tolerance*$_i$ indicates if the household chooses a higher risk tolerance score than that recommended by the robo advisor. *Middle*$_i$ indicates if i belongs to the second ($1k-$6k) or third U.S. wealth quintile ($6k-$42k). Wealth consists of liquid assets, and wealth quintiles are calculated using the Survey of Consumer Finances (SCF) dataset. The sample consists of households who participate with the robo advisor and make a deposit over the period from December 2014 through February 2016. The upper panel summarizes all households in the sample, and lower panel summarizes households from the second or third U.S. wealth quintile. Appendix A contains additional variable descriptions.

Table 2: Democratization of the Robo Market after the Reduction

Outcome:	*New Participant*$_i$					
	(1)	(2)	(3)	(4)	(5)	(6)
Middle$_i$	0.219	0.154	0.141	0.101	0.159	0.166
	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
Controls						
Age$_i$		0.004	0.003	0.001	0.003	0.004
		(0.000)	(0.000)	(0.038)	(0.000)	(0.000)
$\log(Income_i)$		-0.128	-0.111	-0.133	-0.108	-0.116
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
High Risk Tolerance$_i$		-0.030	-0.027	-0.025	-0.025	-0.025
		(0.036)	(0.059)	(0.083)	(0.084)	(0.167)
Measure of Middle	Second or Third Quintile			Second Quintile	Middle with Buffer	Imputed Middle
State FE	No	No	Yes	Yes	Yes	Yes
R-squared	0.033	0.065	0.094	0.085	0.096	0.107
Number of Observations	9,349	9,349	9,349	9,349	8,982	5,728

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

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

where $Middle_i$ indicates if i belongs to the second (\$1k-\$6k) or third U.S. wealth quintile (\$6k-\$42k), as opposed to the fourth or fifth quintile (>\$42k) that together constitute the reference group; and $New\ Participant_i$ indicates if i becomes a robo participant after the reduction, as opposed to before it. Columns (4)-(6) assess the scope for measurement error by remeasuring $Middle_i$ using alternative measures: an indicator for whether i belongs to the second U.S. wealth quintile (Second Wealth Quintile), which we later denote in Table 6 as $Lower\ Middle_i$; an indicator for whether i belongs to the second or third U.S. wealth quintile, after assigning a missing value to households whose liquid assets are within a 10% buffer of the third U.S. wealth quintile (Middle with Buffer); and an indicator for whether i belongs to the second or third U.S. wealth quintile, after assigning a missing value to households whose reported wealth quintile does not equal their imputed wealth quintile based on the boosted trees imputation methodology described in Section 6.1 applied to the SCF dataset. The sample consists all robo participants in the Wealthfront dataset. Wealth consists of liquid assets, defined in Table 1. Wealth quintiles are calculated using the Survey of Consumer Finances (SCF) dataset. Household controls are $\log(Income_i)$, Age_i, and $High\ Risk\ Tolerance_i$, defined in Table 1. Standard errors are clustered by household.

Table 3: Magnitude of Effect on Robo Participation

	Growth in Number of Robo Participants					
	All Participants			Middle-Class Participants		
	Observed (g) (1)	Counterfactual (g^C) (2)	Effect (η) (3)	Observed (g) (4)	Counterfactual (g^C) (5)	Effect (η) (6)
Baseline Estimates:						
Table 2, Column (3)	119.4%	105.7%	13.7%	239.4%	129.8%	109.6%
Additional Estimates:						
Table 2, Column (4)	119.4%	118.4%	1%	165.3%	109.3%	55.9%
Table 2, Column (5)	119.4%	105.3%	14.1%	256.2%	127.6%	128.6%
Table 2, Column (6)	119.4%	104.0%	15.4%	242.2%	118.2%	124.0%

Note: This table summarizes the observed and counterfactual growth rates in the number of robo participants around the reduction, which assesses the magnitude of the results in Table 2. Column (1) summarizes the observed growth rate in the total number of robo participants, denoted g, and column (2) summarizes the counterfactual growth rate in the absence of the reduction, denoted g^C and defined in equation (7) as

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

Column (3) summarizes the effect of the reduction, defined as the difference between g and g^C, that is, $\eta = g - g^C$. Columns (4)-(5) summarize the analogous observed and counterfactual growth rates in the number of middle-class robo participants, and column (6) summarizes the analogous value of η. The counterfactual growth rate in the number of middle-class robo participants in column (5) is defined as

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

Each row calculates these statistics using the estimated β and definition of $Middle_i$ from the indicated specification in Table 2. The observed growth rate in the number of middle-class participants differs across specifications in column (4) because the definition of $Middle_i$ varies across specifications. The remaining notes are the same as in Table 2.

Table 4: Robustness of Investment Constraints as the Channel

	Constrained Investment Measures		Financial Inclusion Measures			
Outcome:	*Under Minimum$_i$*	*At Minimum$_i$*	*Asset Management$_i$*	*Stock Market Participant$_i$*	*Homeowner$_i$*	*High Dividend Zip Code$_i$*
	(1)	(2)	(3)	(4)	(5)	(6)
Middle$_i$	0.295	0.253	-0.180	-0.429	-0.152	-0.056
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Middle$_i$ $\times$ *New Participant$_i$*		-0.317				
		(0.000)				
New Participant$_i$		-0.149				
		(0.000)				
Controls	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.155	0.022	0.213	0.328	0.406	0.540
Number of Observations	5,088	6,890	5,088	5,088	5,088	4,998

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

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

where Y_i is a measure of constraints imposed on i by the previous account minimum. Columns (1)-(2) consider measures based on the household's investment behavior: *Under Minimum$_i$* indicates if i's initial deposit is less than the previous account minimum ($5k); and *At Minimum$_i$* indicates if i's initial deposit equals the previous account minimum or is no more than 5% higher. Column (2) tests for a change in bunching behavior by middle-class participants by including the interaction between *Middle$_i$* and *New Participant$_i$*. Columns (3)-(6) consider measures based on the household's financial inclusion before the reduction: *Asset Management$_i$* indicates if i is imputed to have participated in asset management before the reduction, based on the boosted trees imputation described in Section 6.1 applied to the SCF dataset; *Stock Market Participant$_i$* indicates if i is imputed to have participated in the stock market before the reduction; *Homeowner$_i$* indicates if i is imputed be a homeowner; and *High Dividend Zip Code$_i$* indicates if the zip code associated with i's income bracket and state of residence has an above-median share of households reporting dividend income, based on the IRS SOI Tax Stats dataset. Lower values of the financial inclusion measures proxy for greater constraints imposed by the previous account minimum. The sample consists of households who become robo participants after the reduction, except in column (2) where households who became robo participants before the reduction are also included. The remaining notes are the same as in Table 2.

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

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

Note: P-values are in parentheses. This table estimates equation (9), which assesses the robustness of the baseline results to a dynamic specification that accounts for various time-varying factors that may disproportionately affect middle-class households. Subscripts i and t index household and week. The regression equation in columns 1-3 is of the form

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

where $Post_t$ indicates if t is greater than the week of the reduction; and $New\ Participant_{i,t}$ indexes if i becomes a robo participant in week t, as opposed to the other weeks in our observation window. Columns (2)-(3) include the interaction between $Middle_i$ and a measure of the robo advisor's visibility: $Monthly\ News\ Articles_t$ is the number of news articles about the robo advisor published in the month of week t, a proxy for media attention; and $Monthly\ Advisor\ Blogs_t$ is the number of blog posts written by the robo advisor in the month of week t, a proxy for advertising. Columns (4)-(5) replace $Post_t$ with an indicator for whether t is k months before or after the reduction, respectively denoted $Months\ Before_{t,k}$ and $Months\ After_{t,k}$, where the reference group consists of the month before the reduction ($Months\ Before_{t,1}$). Column 5 includes the interaction between $Post_t$ and an indicator for whether i is under 35 years old, denoted $Millennial_i$. The set of news articles used to construct $News\ Articles_t$ are the top 150 articles, sorted by relevance, from a Google News search of the advisor's name ("Wealthfront") among articles published in 2015. Standard errors are two-way clustered by household and week. The remaining notes are the same as in Table 2.

Table 6: Implications for Inequality in Total Portfolio Returns

Outcome:		$\Delta Risky\ Share_i$				$\Delta Total\ Return_i$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Middle_i$	0.145	0.132	0.133		1.223	1.044	1.056	
	(0.000)	(0.000)	(0.000)		(0.000)	(0.000)	(0.000)	
$Lower\ Middle_i$				0.262				2.084
				(0.000)				(0.000)
$Upper\ Middle_i$				0.121				0.959
				(0.000)				(0.000)
Household Controls	No	Yes	Yes	Yes	No	Yes	Yes	Yes
State FE	No	No	Yes	Yes	No	No	Yes	Yes
R-squared	0.076	0.078	0.092	0.100	0.102	0.111	0.122	0.133
Number of Observations	4,679	4,679	4,679	4,679	4,668	4,668	4,668	4,668

Note: P-values are in parentheses. This table estimates variants of equation (2), which assess the effect of reducing the account minimum on new robo participants' risky share and total portfolio return. Subscript i indexes household. The regression equation is of the form

$$\Delta Y_i = \beta Middle_i + \delta X_i + \tau + u_i.$$

The outcomes ΔY_i are: the ratio of robo investment over the post-reduction period to the household's liquid assets, denoted $\Delta Risky\ Share_i$; and the change in total portfolio return, denoted $\Delta Total\ Return_i$ and defined in equation (12) as

$$\Delta Total\ Return_i \equiv \Delta Risky\ Share_i \times Risky\ Return_i,$$

where $Risky\ Return_i$ is a measure of the expected return on the robo portfolio. The baseline measure, used in this table, is the expected return implied by the CAPM. Columns (4) and (8) decompose $Middle_i$ into households from the second U.S. wealth quintile ($1k-$6k), denoted $Lower\ Middle_i$, and the third quintile ($6k-$42k), denoted $Upper\ Middle_i$. Returns are excess of the risk-free return and net of all fees, including the advisor's management fee. The sample consists of households who become robo participants after the reduction. The remaining notes are the same as in Table 2.

Table 7: Robustness of Implications for Inequality in Returns

Outcome:	$\Delta\,Total\ Return_i$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Middle_i$	1.056	1.168	1.273	0.999	-0.080	1.066	1.071	1.047
	(0.000)	(0.000)	(0.000)	(0.000)	(0.763)	(0.000)	(0.000)	(0.000)
Measure of Return	CAPM	Fama-French	Fama-French with Bond	3-Year Realized	CAPM	CAPM	CAPM	CAPM
Participants in Sample	New	New	New	New	Existing (Placebo)	Stock Market Nonparticipant	No Outflow	Retirement
Household Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.122	0.120	0.106	0.113	0.007	0.107	0.122	0.158
Number of Observations	4,668	4,668	4,668	4,679	3,721	1,698	4,481	566

Note: P-values are in parentheses. This table assesses the robustness of the effect of reducing the account minimum on new robo participants' total portfolio return. Subscript i indexes household. The specification is the same as in columns (5)-(7) of Table 6. Columns (1)-(4) assess robustness to the measure of expected return on the robo portfolio: the expected return implied by the CAPM (CAPM); the expected return implied by the Fama-French Three Factor Model (Fama-French); the expected return implied by the Fama-French Three Factor Model augmented with U.S. and global bond returns (Fama-French with Bond); and the portfolio's realized return over the 3-year period ending in 2018 (3-Year Realized). The sample in columns (1)-(4) consists of households who become robo participants after the reduction (New). Columns (5)-(8) assess robustness to various other subsamples: households who participated with the robo advisor before the reduction (Existing), which serves as a placebo sample; new robo participants who are imputed to be first-time stock market participants, based on the boosted trees imputation described in Section 6.1 applied to the SCF dataset (Stock Market Nonparticipant); new robo participants who do not withdraw any money over the post-reduction period (No Outflow), defined as July 7, 2015 to February 29, 2016; and new robo participants who open a nontaxable, retirement portfolio (Retirement). The remaining notes are the same as in Table 6.

Table 8: Welfare Implications for New Middle-Class Robo Participants

	Median across Robo Participants	
	New Middle-Class (1)	New Upper-Class (2)
Excessive Risky Share$_i$		
Baseline Risk Aversion ($\gamma = 5$)	-0.1	-0.22
Higher Risk Aversion ($\gamma = 7$)	0	-0.12
Compensated Risk Share$_i$		
CAPM	0.64	0.63
Fama-French	0.95	0.95
Fama-French with Bond	0.96	0.96
Probability of Persistence$_i$		
No Account Closure	0.97	0.98
No Subsequent Outflow	0.94	0.96
No Large Subsequent Outflow	0.98	0.98
Subsequent Inflow	0.70	0.71

Note: This table summarizes various measures related to the welfare of households who become robo participants after the reduction. Subscript i indexes household. The upper panel summarizes *Excessive Risky Share$_i$*, defined in equation (15) as the difference between the household's actual risky share and the risky share recommended by the Merton (1969) formula, based on coefficients of relative risk aversion (γ) of 5 or 7. The risky share for former nonparticipants in the stock market equals Δ*Risky Share$_i$*, defined in equation (11) as the ratio of robo investment over the post-reduction period to the household's liquid assets. The risky share for stock market participants cannot be calculated because *Risky Share$_{i,0}$* is unobserved in equation (10). Stock market participation is based on the boosted trees imputation described in Section 6.1 applied to the SCF dataset. The middle panel summarizes *Compensated Risk Share$_i$*, defined as the share of variance in the household's robo portfolio that is compensated according to the asset pricing model used to calculate *Risky Return$_i$*. The three models used are: the CAPM; the Fama-French Three Factor Model (Fama-French); and the Fama-French Three Factor Model augmented with U.S. and global bond returns (Fama-French with Bond). The lower panel summarizes *Probability of Persistence$_i$*, defined as the probability of exhibiting persistent investment behavior and calculated as the share of new robo participants who, over the post-reduction period (July 7, 2015 to February 29, 2016): do not close their account (No Account Closure); do not withdraw any money (No Subsequent Outflow); do not withdraw more than 20% of their initial deposit (No Large Subsequent Outflow); or make a subsequent deposit (Subsequent Inflow). The first column summarizes new robo participants from the second and third quintiles of the U.S. wealth distribution ($1k-$42k), and the second column summarizes new robo participants from the fourth and fifth quintiles (>$42k). The remaining notes are the same as in Table 2.

Online Appendix

This document contains additional material referenced in the text. Appendix A builds on Section 4 by describing our data in greater detail. Appendix B characterizes a stylized model grounded in the framework sketched in Section 2. Appendix C describes how we impute variables not observed in our robo advising dataset. Appendix D performs an aggregated analysis referenced in Section 6.4. Appendix E describes how we estimate expected risky return as used in Section 7.2.

A Additional Data Description

We provide additional details on the paper's two principal datasets: a weekly panel of deposit activity by robo participants (A.1); and a cross-section of all U.S. households (A.2). We also provide a catalog of the paper's key variables (A.3).

A.1 Robo Advising Dataset

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

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

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

We cross-referenced our robo advising dataset against publicly available SEC ADV filings. According to these filings, Wealthfront reported 18,800 participants (i.e., clients) in December 2014 and 61,000 participants in February 2016. As described in the text, the discrepancy between the SEC ADV filings and our dataset is explained by the SEC's filing requirements. Specifically, the SEC states: "The definition of 'client' for Form ADV states that advisors must count clients who do not compensate the advisor" (SEC 2017). Thus, the number of participants reported to the SEC by Wealthfront or any other robo advisor includes participants who did not make any deposits over the sample period as well as "participants" who created a username but never funded a Wealthfront account.

A.2 Survey of Consumer Finances (SCF)

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

We calculate two sets of variables using the SCF dataset. First, we calculate quintiles of the overall U.S. distribution of liquid assets. To maximize comparability with our robo advising dataset, we define liquid assets in the SCF as the sum of checking accounts, savings accounts, certificates of deposit, cash, stocks, bonds, savings bonds, mutual funds, annuities, trusts, IRAs, and employer-provided retirement plans. This definition of liquid assets most closely matches the definition in our robo advising dataset, although the two are not equivalent. For example, we include bonds and savings bonds in the SCF definition, although they are not explicitly mentioned as a liquid asset in the robo advisor's questionnaire. Removing bonds and savings bonds from the SCF definition has little impact because it only changes the boundary between the middle and upper classes by 1%. We carefully examine how measurement error might affect our results in Sections 5.3.1 and 7.1 of the text. Table A3 reports the boundaries that define the five quintiles.

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

A.3 Description of Variables

- *Liquid Assets*$_i$: This variable is the sum of cash, savings accounts, certificates of deposit, mutual funds, IRAs, 401ks, and public stocks for household i, based on the robo advising dataset.

- *Middle*$_i$: This variable indicates if household i's liquid assets fall within the second or third U.S. quintile of liquid assets. Household i's liquid assets are calculated using the robo advising dataset. Quintiles of liquid assets are calculated using the SCF dataset.

- *New Participant*$_i$: This variable indicates if household i becomes a participant with the robo advisor over the period from July 7, 2015 through February 29, 2016. Explicitly, it equals 1 for such individuals and equals 0 for households who participated before July 7, 2015.

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

- *Income*$_i$: This variable is annual income for household i, based on the robo advising dataset.

- *Age*$_i$: This variable is the age of the householder for household i, based on the robo advising dataset.

- *High Risk Tolerance*$_i$: This variable indicates if household i chooses a higher risk tolerance score than recommended by the robo advisor, based on the robo advising dataset.

- *Under Minimum*$_i$: This variable indicates if household i's initial deposit with the robo advisor is less than \$5,000, based on the robo advising dataset.

- *At Minimum*$_i$: This variable indicates if household i's initial deposit with the robo advisor is between \$5,000 and \$5,250, based on the robo advising dataset.

- *Asset Management*$_i$: This variable indicates if household i is imputed to participate in asset management. The imputation is based on the boosted trees algorithm described in Appendix C. This algorithm is trained using the SCF dataset and applied to the robo advising dataset. Participation in asset management is defined based on the SCF dataset as both: (a) owning stocks, mutual funds, a trust, or an IRA; and (b) consulting with a broker, financial planner, banker, accountant or lawyer regarding investment.

- *Stock Market Participant*$_i$: This variable indicates if household i is imputed to participate in the stock market. The imputation is based on the boosted trees algorithm described in Appendix C. This algorithm is trained using the SCF dataset and applied to the robo advising dataset. Participation in the stock market is defined as owning stocks, mutual funds, a trust, or an IRA, based on the SCF dataset.

- *Homeowner$_i$*: This variable indicates if household i is imputed to own owner-occupied residential real estate. The imputation is based on the boosted trees algorithm described in Appendix C. This algorithm is trained using the SCF dataset and applied to the robo advising dataset. Ownership of owner-occupied residential real estate is based on the SCF dataset.

- *High Dividend Zip Code$_i$*: This variable indicates if the zip code associated with i's bracket of annual income and state of residence has a share of households reporting dividend income in 2015 that is greater than or equal to the median share. Data on annual income are from the robo advising dataset, and annual income brackets are defined by the following boundaries: \$25,000; \$50,000; \$100,000; and \$125,000. Data on state of residence are from the robo advising dataset. Data on the share of households in a zip code reporting dividend income are from the IRS SOI Tax Stats dataset, which is available at the zip code level. Zip codes are assigned to an annual income bracket based on the zip code's average adjusted gross income across tax returns. For each income bracket by state bin, we calculate the average zip code-level share of households reporting dividend income. Then, we calculate the median of this share across bins.

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

- Δ*Total Return$_i$*: This variable is the product of Δ*Risky Share$_i$* and *Risky Return$_i$*, defined as a measure of the expected return on the portfolio that the robo advisor manages for household i. The baseline measure of expected return is the expected return implied by the CAPM. Appendix E describes additional measures. Robo portfolios are indexed by risk tolerance score and tax status. Data on Δ*Risky Share$_i$* are from the robo advising dataset. Data used to calculate *Risky Return$_i$* for a given portfolio are from the Center for Research in Security Prices (CRSP) and Ken French's website.

B Stylized Model

We propose a stylized model that exemplifies the theory sketched in Section 2. As in Section 2, households solve a portfolio optimization problem in which they delegate a share ρ of their risky portfolio to asset managers. We take the limiting case as ρ approaches one. Asset managers offer a portfolio with risky, net-of-fee return R. Alternatively, households can invest in a riskless asset that delivers return R^f, where, to minimize notation, $R^f = 0$.

Asset managers have a cost structure such that they require an account minimum M. Therefore, given investable assets W, the household solves the following version of the traditional mean-variance optimization problem,

$$\max_{\omega} \quad U(\omega) = \left\{ \omega \mathbb{E}\left[R\right] - \frac{1}{2}\gamma\omega^2 \mathrm{Var}\left[R\right] \right\} \tag{B1}$$

$$s.t.$$

$$\omega W \geq M, \tag{B2}$$

where ω denotes the share of investable assets allocated to the advisor. One can motivate the problem in (B1) by supposing R is normally distributed and uncorrelated with nonfinancial income, and the household has a constant coefficient of absolute risk aversion γ, not to be confused with the use of γ to denote the coefficient of relative risk aversion in the text.

Absent the constraint in (B2), households invest in the risky portfolio and allocate a share

$$\tilde{\omega} \equiv \frac{1}{\gamma} \frac{\mathbb{E}\left[R\right]}{\mathrm{Var}\left[R\right]} \tag{B3}$$

of their investable wealth W to the robo advisor. Under the account minimum M, this optimal allocation may no longer be feasible, and so the household's participation status depends on wealth. In particular, the solution to problem (B1) is

$$\omega^* = \begin{cases} 0, & W \leq \max\left\{M, \frac{M}{2\tilde{\omega}}\right\} \\ \frac{M}{W}, & \max\left\{M, \frac{M}{2\tilde{\omega}}\right\} < W \leq \frac{M}{\tilde{\omega}} \\ \tilde{\omega}, & \frac{M}{\tilde{\omega}} < W \end{cases}, \tag{B4}$$

and the household's participation status is given by

$$p^* = \begin{cases} 0, & W \leq \max\left\{M, \frac{M}{2\tilde{\omega}}\right\} \\ 1, & \max\left\{M, \frac{M}{2\tilde{\omega}}\right\} < W \end{cases}. \tag{B5}$$

The wealth cutoff for participation in equation (B5) does not necessarily equal the account minimum, M, and it is likely to be larger than M if $\tilde{\omega}$ is reasonably low. Intuitively, if the household's wealth lies below the cutoff, it would either need to invest such a large fraction of its wealth in the risky portfolio that it optimally chooses not to participate, in which case the cutoff equals $\frac{M}{2\tilde{\omega}}$, or it simply does not have enough assets to overcome the minimum without borrowing, in which case the cutoff equals M.

Now consider a reduction in the account minimum from M to M'. It is straightforward to show that the household responds by becoming a participant if and only if

$$\underline{W} \equiv \max\left\{M', \frac{M'}{2\tilde{\omega}}\right\} < W \leq \max\left\{M, \frac{M}{2\tilde{\omega}}\right\} \equiv \overline{W}. \tag{B6}$$

Therefore, the framework predicts that the effects of the reduction will be concentrated among households with intermediate levels of wealth (i.e., "middle class"). Households with low levels of wealth such that $W < \underline{W}$ (i.e., "lower class") are predicted to remain nonparticipants, while households with high levels of wealth such that $W > \overline{W}$ (i.e., "upper class") remain participants. As discussed in Section 4.2 of the text, the empirical analogues to the lower, middle, and upper classes are households from the first, second or third, and fourth or fifth quintiles of the U.S. distribution of liquid assets, respectively.

Turning to the wealth distribution across participants, let ΔP_m and ΔP_u respectively denote the changes in the total number of middle-class and upper-class participants caused by the reduction. According to the framework,

$$\Delta P_m = \int_{\underline{W}}^{\overline{W}} \left[p^*(W, M') - p^*(W, M)\right] dW, \tag{B7}$$

$$\Delta P_u = \int_{\overline{W}}^{\infty} \left[p^*(W, M') - p^*(W, M)\right] dW, \tag{B8}$$

which implies the following difference-in-difference equation

$$\Delta P_m - \Delta P_u = \int_{\underline{W}}^{\overline{W}} \left[p^*(W, M') - p^*(W, M)\right] dW > 0. \tag{B9}$$

In words, equation (B9) says that the predicted increase in participation by the middle class exceeds that of the upper class. The results in Tables 2, 5, and A4 all support this prediction.

Letting $\bar{R} \equiv \omega \mathbb{E}\left[R\right]$ denote a household's total portfolio return, the difference-in-difference between a new middle-class participant's total portfolio return, $\bar{R}_m$, and a new upper-class participant's total portfolio return, $\bar{R}_u$, is

$$\Delta\bar{R}_m - \Delta\bar{R}_u = \Delta\omega_m^* \mathbb{E}\left[R\right] > 0, \tag{B10}$$

where

$$\Delta\omega_m^* = \begin{cases} \frac{M'}{W}, & \max\left\{M', \frac{M'}{2\tilde{\omega}}\right\} < W \leq \frac{M'}{\tilde{\omega}} \\ \\ \tilde{\omega}, & \frac{M'}{\tilde{\omega}} < W \leq \max\left\{M, \frac{M}{2\tilde{\omega}}\right\} \end{cases}. \tag{B11}$$

Equation (B10) shows how inequality in total return falls as middle-class households experience an increase in total return. As shown in Table 8, this increase in total return primarily reflects greater compensated risk.

Relative to the discussion in Section 2 of the text, equation (B9) summarizes our prediction of "asymmetric diversification" using a closed-form expression. We find strong support for this prediction in Figure 2 and Table 2. Moreover, Table 6 supports the prediction from equation (B10) that inequality in total return falls. Lastly, our model implies that, by revealed preference, middle-class households benefit from this increase in total return. In reality, households may take excessive risk, not diversify away uncompensated risk, withdraw their investment prematurely, or make other mistakes that cannot be captured by our stylized model. While we cannot rule out any of these possibilities, the analysis in Section 8 suggests that middle-class households with conventional utility functions do indeed benefit from automated asset management.

C Imputation Methodology

As described in Section 6.1.3, we use the SCF dataset impute various measures of financial inclusion for households in our robo advising dataset, namely, a household's participation in the stock market, professional asset management, and homeownership. Imputation of these variables is necessary because we do not observe them in our robo advising dataset. We also impute a household's wealth class to assess the scope for bias due to measurement error in Section 5.3.1. Our imputation procedure involves predicting a variable's value for each household in our robo advising dataset using a prediction model applied to the set of variables observed in both our robo advising dataset and the SCF dataset. Accordingly, we first estimate, cross-validate, and test each model on the SCF dataset. A priori, we take no stance on which model to use. Instead, we let the data inform which model to use, based on performance in out-of-sample tests.

In this appendix, we describe the candidate prediction models (C.1), the process for imputing measures of financial inclusion (C.2), and the process for imputing a household's wealth class (C.3)

C.1 Prediction Models

We consider four parametric prediction models: standard logistic regression, lasso logistic regression, ridge logistic regression, and elastic net logistic regression. The latter three models all involve regularizations applied to the standard logistic regression model, and so we collectively refer to the lasso, ridge, and elastic net logistic regression models as logistic regression with regularizations. We also consider two nonparametric, tree-based models: random forest and boosted trees. We now describe each model and its hyperparameters, that is, a set of model-specific parameters that govern how the estimation is performed but have no economic interpretation themselves.

C.1.1 Standard Logistic Regression Model

Logistic regression models assume that the relationship between the dependent variable and independent variables takes a particular functional form. To estimate a standard logistic regression model, we minimize the following negative log-likelihood function across the sample of N observations:

$$L(y, \zeta, \psi, x) = -\sum_{i=1}^{N} \left(y_i \log\left(\xi_i\right) + (1 - y_i) \log\left[1 - \xi_i\right] \right), \tag{C1}$$

with

$$\xi_i \equiv \frac{1}{1 + \exp(-(x_i \psi + \zeta))}, \tag{C2}$$

where $y_i \in \{0, 1\}$; the model parameters are given by $\psi = (\psi_1, \ldots, \psi_q)'$ and ζ; and the set of predictors is given by $x = (x_1, \ldots, x_q)$.

C.1.2 Logistic Regression Models with Regularizations

Ridge, lasso, and elastic net are all classic shrinkage methods employing techniques that optimally shrink the coefficient estimates. We augment the baseline logistic regression with three regularizations to improve the model's performance.

First, we introduce a ridge regularization. The objective function to be minimized is

$$\kappa \sum_{j=1}^{q} \psi_j^2 + L(y, \zeta, \psi, x). \tag{C3}$$

The penalty term is $\kappa \sum_{j=1}^{q} \psi_j^2$, which is called the l_2 penalty, and $\kappa > 0$ is a hyperparameter that controls the model's complexity. By increasing κ, the level of shrinkage becomes higher, which limits on the size of the coefficients. The predictors need to be standardized before running a ridge regression. When predictors are highly correlated with each other (e.g., wealth and income), the ridge penalty enforces similarity in their coefficient estimates.

Second, we introduce a lasso regularization, which constrains the absolute values of the estimated coefficients. The objective function for lasso is

$$\kappa \sum_{j=1}^{q} |\psi_j| + L(y, \zeta, \psi, x). \tag{C4}$$

The penalty term in lasso regression is $\kappa \sum_{j=1}^{q} |\psi_j|$, which is named the l_1 penalty, and κ is the hyperparameter. This penalty term could drive some coefficients to zero if κ is large enough. As in ridge regression, we need to standardize the independent variables before running the regression. The main difference between ridge and lasso regressions is that coefficient estimates in ridge regression rarely equal zero exactly, while lasso regression coefficient estimates are more likely to equal zero. This feature gives lasso regression a natural advantage when the true model is sparse and there are many redundant predictors. The limitation of lasso is that it may only select one variable out of a group of highly correlated variables.

Third, we introduce an elastic net regularization. Under this regularization, the estimated coefficients are values that minimize the objective function

$$\kappa \sum_{j=1}^{q} \left(\varrho \psi_j^2 + (1 - \varrho)|\psi_j| \right) + L(y, \zeta, \psi, x). \tag{C5}$$

The penalty term for elastic net is $\kappa \sum_{j=1}^{q} (\varrho \psi_j^2 + (1 - \varrho)|\psi_j|)$ and there are two hyperparameters, κ and ϱ. The elastic net penalty is a linear combination of the lasso and the ridge penalties: it is able to select a subset of predictors as in the lasso, and it drives the coefficients of correlated predictors down simultaneously as in the ridge.

C.1.3 Nonparametric Models

We next describe nonparametric, tree-based models. While linear regularization-based models could perform better than the baseline logistic regression, especially when there are many correlated variables, these models still impose a parametric relationship between the predictors and the target variable. Random forest and boosted trees are both tree-based models, which allow for the inclusion of complex interactions among the predictors.

The family of tree-based models divides the space of predictors into an assortment of rectangles S_j and then fits a simple model (e.g., a constant ϕ_j) within each rectangle partition. The formal expression of a tree is

$$T(x;\Theta) = \sum_{j=1}^{J} \phi_j \mathbb{1}\left[x \in S_j\right] \tag{C6}$$

where $\Theta = \{S_j, \phi_j\}_1^J$ encodes how the sample is partitioned into rectangles.

To make equation (C6) more concrete, consider an example of predicting a household's stock market participation status using the household's age and income as independent variables, so that $x = (Age, Income)$. In this example, one such rectangle S_j may be the subsample of households with age between the boundaries $\underline{Age}$ and $\overline{Age}$ and annual income between the boundaries $\underline{Income}$ and $\overline{Income}$. The corresponding value of ϕ_j would be the stock market participation rate among households whose age and income place them within rectangle S_j.

The partition Θ associated with tree T is found by minimizing some loss function

$$\hat{\Theta} = arg\,min_{\Theta} \sum_{j=1}^{J} \sum_{x_i \epsilon S_j} L(S_j). \tag{C7}$$

In our procedure, the loss function $L(S_j)$ is the Gini index within region S_j, defined as

$$L(R_j) = 1 - \left[\hat{\phi}_j^2 + (1 - \hat{\phi}_j)^2\right]. \tag{C8}$$

In the context of the previous example, $\hat{\phi}_j$ would be the stock market participation rate among households in rectangle S_j, based on the subsample of the SCF dataset used to estimate the model. Intuitively, the loss function in equation (C8) rewards rectangles where the share of stock market participants is close to either zero or one.

An optimization routine called recursive binary splitting is used to optimally divide the set of independent variables into space into J rectangles. Again making use of the previous concrete example, this optimization routine would select the values of $\underline{Age}$, $\overline{Age}$, $\underline{Income}$, and $\overline{Income}$ that minimize equation (C8). When there are many independent variables, recursive binary splitting selects not only the cut-points that define each rectangle S_j but also the set of independent variables used to partition the sample. The splitting process ends after a certain number of splits, called, the "tree size". Once the tree is "grown", the predictions are formed based on the prevailing values within the rectangles: if most households from the training set in the rectangle are stock market participants, a household from the test set who falls into the same rectangle is also predicted to be a stock market participant.

In any tree-based model, the number of sample splits (i.e., tree size) is an important hyperparameter. A large tree with many splits may lead to overfitting, while a small tree with relatively few splits may lead to inaccurate predictions. A single-tree model produces approximately unbiased estimates, but these estimates can be inefficient, that is, have a large variance. Two variants of the single-tree model that offer substantially more efficient estimates are a random forest model and a boosted regression trees model. We discuss the intuition behind these models below and defer complete methodological details to James et al. (2013).

The random forest model introduces a technique called “bagging”, or bootstrap aggregation. The bagging process reduces the variance through averaging predictions from many trees. In particular, a random forest is a collection of uncorrelated trees, which is created by growing trees with a certain degree of randomization. The model first draws multiple bootstrap samples from the training data, and fits a tree for each sample with a random subset of the predictors at each split. Specifically, random forests only consider a subset m of all available predictors q at each split, typically $m = \sqrt{q}$. Next, the predictions are obtained by averaging the results across the fitted trees from each sample. The random forest’s hyperparameters are the number of trees and the number of splits for each tree.

The boosted trees method creates a tree ensemble using the idea of boosting. Boosting techniques seek to improve the prediction power by training a sequence of “weak” models, each compensating the weakness of its predecessors. Specifically, a “boosted trees” algorithm grows the trees in a sequential way such that each new tree is constructed using information about the residuals from the previous tree. In the stock market participation example, the boosting algorithm would optimize equation (C8) after disproportionately penalizing prediction errors in the current tree for households whose prediction errors in the previous tree were larger. The degree of penalty is governed by a shrinkage hyperparameter.

There are three key hyperparameters for models that incorporate boosting: the number of trees, the shrinkage parameter, and the number of splits. In general, a boosted trees model features shallow trees with relatively few splits (i.e., “weak”). The number of splits d in each tree controls the model’s complexity. A higher value of d implies more complex interactions among the predictors. The shrinkage parameter controls the speed with which the algorithm learns from previous prediction errors, and a very small shrinkage parameter should be combined with a very large number of trees for better prediction accuracy.

C.1.4 Hyperparameters

Based on the previous model descriptions, each model has a set of hyperparameters that govern how the estimation is performed. Typically, hyperparameter values are chosen to optimize a predetermined performance metric on a subset of the full dataset, called the validation set. This validation set is distinct from the subset used to train the prediction model, called the training set. The final subset of the data, called the testing set, is used to evaluate the out-of-sample performance of the model with the optimally chosen hyperparameter values. The process of optimizing hyperparameter values is called hyperparameter tuning.

We implement a grid search procedure, which involves an exhaustive search through a predetermined grid of hyperparameter values and is the traditional method for hyperparameter optimization (Murphy 2012). To choose among various sets of hyperparameter values, we use a 10-fold cross-validation procedure.

C.2 Imputing Financial Inclusion Measures

We use the following procedure to impute participation in the stock market, professional asset management, and homeownership. To avoid overfitting, we first split the sample following a standard 80/20 sample split. In particular, we use 80% of the data as a training/cross-validation sample to tune each model's hyperparameters, and we test the model using the remaining 20% of the data as a testing sample. All of the models are estimated using *scikit-learn* library in Python.

As mentioned above, we start by estimating each model on the SCF training sample to later impute the measures of financial inclusion that are unobserved in our robo advising dataset. We apply each model to the independent variables that are jointly observed across the two datasets: liquid assets, income, and age. Then, we use a 10-fold cross-validation procedure to choose the hyperparameters. The hyperparameter optimization is based on the receiver operating characteristic (i.e., ROC) curve, and we choose hyperparameters to maximize the area under this curve (AUC).

We first estimate a standard logistic regression model and use it as a baseline to compare with other models. We next turn to logistic regression models with regularizations. For these three models, we use the cross-validation procedure to jointly choose the optimal regularization and the hyperparameters for that regularization. Finally, we separately fit random forest and boosted trees models, using cross-validation to optimize the hyperparameters. As a result, we effectively have a choice among four models: the baseline logistic regression, the optimally chosen logistic regression model with regularization, and the two tree-based models. Across all the models, we report the average ROC-AUC across 10 validation sets as well as the ROC-AUC in the test set. We also report the accuracy in the test set as an additional performance metric.

C.2.1 Results

The results for stock market participation are reported in Table A7. The tree-based models perform better than the parametric models, suggesting that the relationship between predictors and stock market participation status is highly non-linear. In the case of boosted trees, the ROC-AUC equals 99% in the average validation set and 96% in the test set, and the model's accuracy equals 97%. The random forest model exhibits slightly lower performance metrics: a 95% validation set ROC-AUC, an 87% test set ROC-AUC, and an accuracy of 87%. While we find that lasso is the preferred logistic regression model with regularization, its performance metrics are not very different from the standard logistic regression. They are also significantly lower than among the tree-based models. The results on participation in professional asset management are reported in Table A8. The comparable performance metrics suggest that the boosted trees model performs best. Finally, Table A9 reports the results for homeownership status and, again, the boosted trees model performs best.

Since the boosted trees model exhibits slightly higher performance metrics than the random forest model and substantially higher performance than the parametric models for all three financial inclusion measures, we impute these measures for each household in our robo advising dataset using the boosted trees model. We use these imputed values in Tables 4 and 7.

C.3 Imputing Wealth Class

We follow a similar procedure for imputing the wealth class for households in our robo advising dataset. Explicitly, we first predict wealth class (i.e., lower, middle, or upper class) using the two other independent variables observed in both the robo advising and SCF datasets: age and income. Then, we assess the quality of self-reporting by comparing a household's predicted wealth class and the wealth class associated with the household's self-reported liquid assets (i.e., self-reported wealth class). Given that wealth class is non-binary, we modify the baseline imputation procedure in two ways. First, we do not use parametric models and instead rely only on non-parametric models. Additionally, we do not use the loss function based on ROC-AUC and instead rely on prediction accuracy as the main performance metric in our 10-fold cross-validation procedure.

C.3.1 Results

Table A10 reports the results. The boosted trees prediction model suggests that 58% and 62% of robo participants who self-report as middle-class and upper-class, respectively, are predicted to indeed belong to their self-reported wealth class. We obtain similar results from the random forest model. Thus, according to the imputation procedure, roughly 60% of robo participants are "telling the truth". The similar truth-telling rates among middle and upper-class robo participants suggests that the definition of liquid assets in our robo advising dataset is not systematically biased downward relative to the SCF definition. Otherwise, we would observe a much lower truth-telling rate among the middle class, since, in the case of such systematic downward bias, robo participants who self-report as middle-class in our robo advising dataset would be imputed to belong to the upper class. More likely, the fact that truth-telling rates fall short of 100% reflects a combination of prediction error, which, per the bottom row in Table A10, would apply to 15% of the sample, as well as classical measurement error in self-reported liquid assets. As discussed in Section 5.3.1, classical measurement error in middle-class wealth status would bias the estimates toward zero through attenuation.

D Aggregated Measure of Democratization

As mentioned in the text, the baseline regression equation (2) is estimated on the set of eventual robo participants, which is appropriate given our interest in the democratization of the market for automated asset management. In particular, this regression quantifies the share of robo participants who were brought into the market by the reduction. We now estimate the effects of the reduction on a related statistic: the share of eligible middle-class households who participate with the robo advisor.

Our procedure has two steps. In the first step, we partition the U.S. population into cohorts based on observable characteristics, as described below, and calculate the total number of new robo participants within each cohort and week. Using notation similar to that in Appendix B, the empirical probability that a household in cohort c becomes a robo participant in week t is

$$p_{c,t} = \frac{\mathit{New\ Participants}_{c,t}}{\mathit{Eligible\ Participants}_{c}}, \tag{D1}$$

where c and t index cohort and week; $\mathit{New\ Participants}_{c,t}$ is the number of new robo participants from cohort c in week t; and $\mathit{Eligible\ Participants}_{c}$ is the number of households from cohort c eligible to become robo participants, in that they are both aware of the robo advisor and satisfy the demographic requirements needed to participate (e.g., U.S. resident over 18 years old). Similarly, the probability that any eligible household becomes a robo participant over all weeks t in our sample period is

$$P = \sum_{t} \sum_{c} w_c p_{c,t}, \tag{D2}$$

where w_c is the share of all eligible households in cohort c. We assume that the total number of eligible participants from any given cohort is approximately constant over our sample period.

In the second step, we estimate the effect of the reduction on the percent change in the probability of participation for households in cohort c. Taking logs of equation (D1) and incorporating cohort-level controls and fixed effects, we estimate the following regression equation

$$\log\left(\mathit{New\ Participants}_{c,t}\right) = \beta\left(\mathit{Middle}_c \times \mathit{Post}_t\right) + \alpha_c + \tau_t + \lambda X_{c,t} + u_{c,t}, \tag{D3}$$

where Middle_c indicates if c belongs to the second or third U.S. wealth quintile, in contrast to the fourth or fifth quintiles that comprise the reference group; and

$$\alpha_c \equiv \log\left(\mathit{Eligible\ Participants}_c\right) + \tilde{\alpha}_c \tag{D4}$$

is a composite cohort fixed effect. Like with the household-level analogue in equation (9), middle-class cohorts constitute the treated group in equation (D3), since they experience a relaxation of investment constraints because of the reduction. Thus, the parameter β equals the effect of the reduction on the percent change in $p_{c,t}$, as distinct from the percentage point change. The percentage point change in the overall probability of robo participation, P, is

$$\Delta P = \sum_t \sum_c w_c p_{c,t} \left[e^{\beta(Middle_c \times Post_t)} - 1 \right], \tag{D5}$$

which follows from equation (D2).

In terms of implementation, we partition households into cohorts according to a set of sorting variables. Adding new sorting variables reduces intra-cohort heterogeneity such that we can better isolate variation in wealth, but it also introduces noisy idiosyncratic variation. In light of this tradeoff, we construct three separate partitions that progressively increase the number of sorting variables. First, we define cohorts by sorting the U.S. population into a maximum of 10 bins based on decile of liquid assets, where deciles are calculated using the SCF dataset. Second, we introduce a three-way sort into a maximum of 1,000 bins based on deciles of liquid assets, age, and income, such that there are 10 possible groups for each variable. In our final approach, cohorts are defined by sorting the U.S. population into a maximum of 6,250 bins based on state of residence, of which there are 50 groups, and quintiles of liquid assets, age, and income, of which there are 5 groups each.

The results of cohort-level panel regression are reported in Table A4. The first two columns estimate the effect of the reduction after sorting the U.S. population only by liquid assets. The results in column (2) suggest that the reduction in account minimum increases the number of middle-class robo participants by 29%. The estimated effect equals 9% when estimated from the three-way sort based on liquid assets, age and income, as shown in column (4), and it equals 2% after additionally sorting by state of residence, as shown in column (6). The lower point estimates reflect the fact that more granular bins can contain fewer robo participants, and so, to recover the increase in the overall probability of robo participation, we must aggregate the effects across all cohorts belonging to the middle class, according to equation (D5).

Quantitatively, we can back out the percent change in the overall probability of robo participation when the product $w_c p_{c,t}$ is uniform. In that case, the reduction's effect on the overall probability of participation implied by the estimate in column (4) is 109%, or, to match the statement in Section 6.4, the reduction roughly doubles the baseline probability. Depending on the covariance between w_c and $p_{c,t}$, this estimate may either overstate or understate the true effect.

Qualitatively, our cohort-level results imply a similar conclusion as our baseline results in Table 2: the reduction increases robo participation by middle-class households who, as we find in Section 6.1, were constrained by the previous account minimum. In particular, this core conclusion continues to hold even when observational units are cohorts and, thus, not restricted to eventual robo participants.

E Estimating Expected Return

This appendix describes the method for estimating the expected return on robo portfolios in Section 7. As mentioned in Section 7.2, using a historical average to measure expected return is subject to well-known issues related to limited time horizons, and especially so in our setting given the relatively-short history of many ETFs managed by our data provider. Therefore, we follow Calvet, Campbell and Sodini (2007) and propose an asset pricing model to estimate the expected returns for the securities in our sample.

While imposing a model improves the efficiency of expected return estimates relative to directly measuring them from historical returns, it leads to some bias by imposing an imperfect model of the return structure. Since the choice of model is somewhat arbitrary and the degree of bias will depend on the characteristics of the portfolio in question, we estimate expected return across a variety of common asset pricing models, indexed by their vector of factors F.

For each security k (i.e., ETF k), we estimate the following pricing kernel

$$Return_{k,t} = \beta_k^F F_t + \epsilon_{k,t}^F, \tag{E1}$$

where F_t denotes a column vector of pricing factors in month t; β_k^F denotes the respective row vector of loadings; $Return_{k,t}$ denotes the monthly return on security k in excess of the risk-free return, measured by the one-month Treasury yield, and net of expense ratios and other fees; and $\epsilon_{k,t}^F$ is an idiosyncratic, zero-mean shock to security k with standard deviation σ_k^F. We estimate equation (E1) using the longest available time series of monthly returns for each security k and factor vector F dating back to January 1975.

Given the estimated loadings $\hat{\beta}_k^F$ from estimating equation (E1) for model F, it is straightforward to compute the expected return on household i's robo portfolio, denoted $Risky\ Return_i^F$. Explicitly, if there are K securities and N factors, then

$$Risky\ Return_i^F = w_i \boldsymbol{\beta}^F \pi^F, \tag{E2}$$

where w_i is a $1 \times K$ row vector of weights across securities in household i's robo portfolio; $\boldsymbol{\beta}^F$ is a $K \times N$ matrix of factor loadings; and π^F is an $N \times 1$ column vector of factor risk prices.

We estimate equation (E1) for the following asset pricing models,

$$F_t^{CAPM} = [R_t^m]', \tag{E3}$$

$$F_t^{FF} = \left[R_t^m, \quad R_t^{HML}, \quad R_t^{SMB}\right]', \tag{E4}$$

$$F_t^{FF+} = \left[R_t^m, \quad R_t^{HML}, \quad R_t^{SMB}, \quad R_t^{USB}, \quad R_t^{GLB}\right]', \tag{E5}$$

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

In words, equations (E3)-(E5) are: the standard capital asset pricing model (CAPM), the Fama-French Three Factor Model, and a five-factor model augmenting the Fama-French model with U.S. and global bond returns. Our data on monthly returns come from the Center for Research in Security Prices (CRSP) and Ken French's website, as described in Appendix A. We use the sample mean to calibrate the factor risk prices, multiplied by 12 to obtain an approximate annual value. The corresponding risk prices for the models in equations (E5)-(E5) are

$$\pi^{CAPM} = [0.068]', \tag{E6}$$

$$\pi^{FF} = [0.068, \quad 0.036, \quad 0.004]', \tag{E7}$$

$$\pi^{FF+} = [0.068, \quad 0.036, \quad 0.004, \quad 0.062, \quad 0.060]'. \tag{E8}$$

Similarly, we use the sample covariance matrix to calculate the share of risk that is compensated according to a given asset pricing model, studied in Section 8.2. Explicitly, this share equals

$$\textit{Compensated Risk Share}_i^F = \frac{\left(w_i\boldsymbol{\beta}^F\right)\Sigma^F\left(w_i\boldsymbol{\beta}^F\right)'}{\left(w_i\boldsymbol{\beta}^F\right)\Sigma^F\left(w_i\boldsymbol{\beta}^F\right)' + w_i\Sigma_\epsilon^F w_i'}, \tag{E9}$$

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

Additional Figures and Tables

Figure A1: Robustness of Graphical Evidence

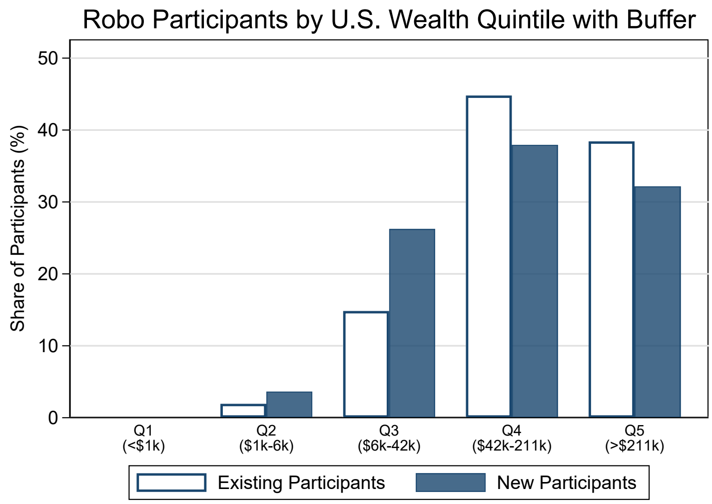

(a) Robo Participants by Buffered Wealth Quintile

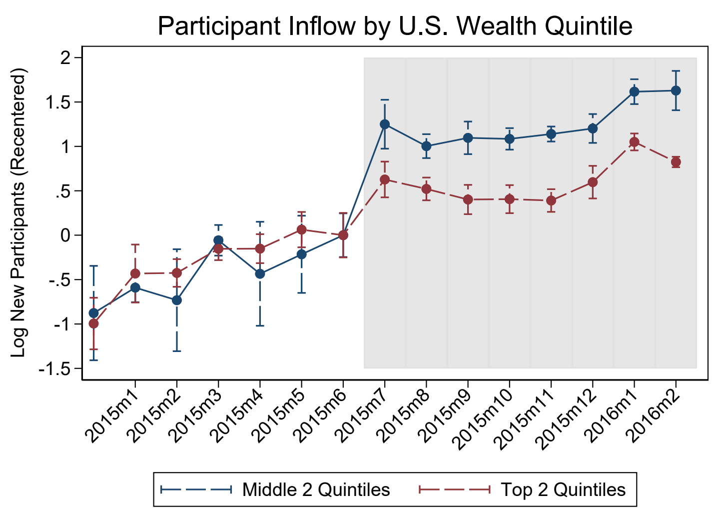

(b) Significance of Pre-Trends in Robo Participation

Note: This figure assesses the robustness of Figures 2 and 3. Panel (a) is analogous to Figure 2 after dropping households whose liquid assets are within a 10% buffer of each U.S. wealth quintile, which assesses robustness to measurement error in liquid assets. Panel (b) is analogous to Figure 3 after demeaning the log number of new participants by its value the month prior to the reduction (June 2015) for each set of wealth quintiles and plotting 95% confidence intervals.

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

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

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

Table A2: Summary of Robo Portfolios

Risk Tolerance (0.5 to 10)	Beta	Expected Return (%)	Stocks (%)	Bonds (%)	Other (%)	Percent of Households (%)
0.50	0.32	4.15	33.00	60.00	7.00	0.67
2.00	0.45	5.15	47.00	48.00	5.00	0.39
2.50	0.49	5.40	50.00	44.00	6.00	0.20
3.00	0.52	5.55	53.00	41.00	6.00	0.89
3.50	0.57	5.95	59.00	35.00	6.00	0.86
4.00	0.58	6.02	59.00	35.00	6.00	1.56
4.50	0.61	6.28	62.00	33.00	5.00	1.14
5.00	0.64	6.51	66.00	29.00	5.00	1.68
5.50	0.67	6.73	69.00	26.00	5.00	1.21
6.00	0.70	6.94	72.00	23.00	5.00	2.27
6.50	0.72	7.10	74.00	21.00	5.00	2.32
7.00	0.75	7.26	77.00	18.00	5.00	6.41
7.50	0.77	7.44	80.00	15.00	5.00	8.06
8.00	0.79	7.61	82.00	13.00	5.00	14.39
8.50	0.82	7.84	86.00	9.00	5.00	16.50
9.00	0.85	8.03	89.00	6.00	5.00	16.30
9.50	0.88	8.26	90.00	5.00	5.00	5.43
10.00	0.91	8.45	90.00	5.00	5.00	19.72

Note: This table summarizes robo portfolios assigned to households in our sample. Portfolios are indexed by risk tolerance score, which ranges from 0.5 to 10 in increments of 0.5, and tax status. Each portfolio has a unique vector of weights across 10 possible ETFs, which are chosen to represent exposure to different asset classes. Stocks, Bonds, and Other respectively denote the sum of weights for ETFs that track stock market indices (VIG, VTI, VEA, VW), bond market indices (LQD, EMB, MUB, SCHP), and other asset classes, namely real estate (VNQ) and commodities (XLE). Expected Return and Beta are based on the CAPM, as described in Section 7.2. The rightmost column shows the percent of robo participants with the indicated portfolio. The table only shows taxable portfolios to emphasize how the allocation varies across risk scores, rather than tax status.

Table A3: Summary of U.S. Wealth Quintiles

Wealth Quintile:	First	Second	Third	Fourth	Fifth
Participation in the Stock Market	0.003	0.064	0.314	0.579	0.87
Participation in Asset Management	0.002	0.041	0.207	0.418	0.693
Conditional Participation in Asset Management	0.66	0.64	0.66	0.72	0.8
Range of Liquid Assets ($000)	[0, 0.6]	[0.6, 6.3]	[6.3, 42]	[42, 211]	> 211

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

Table A4: Robustness to Aggregated Measure of Democratization

Outcome:	$\log\left(New\ Participants_{c,t}\right)$					
	(1)	(2)	(3)	(4)	(5)	(6)
$Middle_c \times Post_t$	0.603	0.290	0.092	0.092	0.022	0.021
	(0.007)	(0.024)	(0.002)	(0.002)	(0.047)	(0.047)
	Wealth Cohorts		Wealth by Income by Age Cohorts		Wealth by Income by Age by State Cohorts	
Controls	No	Yes	No	Yes	No	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.868	0.875	0.390	0.390	0.324	0.324
Number of Observations	455	455	22,013	22,013	52,025	52,025

Note: P-values are in parentheses. This table estimates equation (D3), which assesses the robustness of the baseline results to a specification in which observational units are aggregates and, thus, not limited to eventual robo participants. Subscripts c and t index cohort and week. The regression equation is of the form

$$\log\left(New\ Participants_{c,t}\right) = \beta\left(Middle_c \times Post_t\right) + \alpha_c + \tau_t + \lambda X_{c,t} + u_{c,t},$$

where $Middle_c$ indicates if c belongs to the second or third U.S. wealth quintile (\$1k-\$42k), as opposed to the fourth or fifth quintile (>\$42k); $Post_t$ indicates if t is greater than the week of the reduction; and $New\ Participants_{c,t}$ is one plus the number of new robo participants from cohort c in week t. In columns (1)-(2), cohorts are defined by sorting the U.S. population into a maximum of 10 bins based on decile of liquid assets (10), where deciles are calculated using the SCF dataset. In columns (3)-(4), cohorts are defined by sorting the U.S. population into a maximum of 1,000 bins based on decile of liquid assets (10), age (10), and income (10), where deciles are calculated using the SCF dataset. In columns (5)-(6), cohorts are defined by sorting the U.S. population into a maximum of 6,250 bins based on state of residence (50) and quintile of liquid assets (5), age (5), and income (5), where quintiles are calculated using the SCF dataset. Controls are averages of the household-level controls from Table 2 across robo participants in cohort c and week t. Standard errors are two-way clustered by cohort and week. The remaining notes are the same as in Table 2.

Table A5: Robustness to Measurement Error from Increases in the Savings Rate

Outcome:	Δ*Savings Rate*$_i$		
	(1)	(2)	(3)
Middle$_i$	-0.163	-0.201	-0.198
	(0.000)	(0.000)	(0.000)
Household Controls	No	Yes	Yes
State FE	No	No	Yes
R-squared	0.054	0.123	0.138
Number of Observations	5,108	5,108	5,108

Note: P-values are in parentheses. This table estimates a variant of equation (2), which assesses whether measurement error in risky share biases the results in Table 6 upward because middle-class households increase their savings rate. The specification is the same as in Table 6 with a different outcome variable: Δ*Savings Rate*$_i$ is the ratio of annualized robo investment over the post-reduction period to the household's annual income. The remaining notes are the same as in Table 6.

Table A6: Consistency with Models of Portfolio Choice

Outcome:	Δ *Risky Share*$_i$ (1)	Δ *Total Return*$_i$ (2)
Middle$_i$	0.128	1.023
	(0.000)	(0.000)
Middle$_i$ × *High Risk Tolerance*$_i$	0.043	0.326
	(0.041)	(0.048)
High Risk Tolerance$_i$	-0.010	-0.073
	(0.367)	(0.383)
R-squared	0.067	0.076
Number of Observations	5,088	4,791

Note: P-values are in parentheses. This table estimates a variant of equation (2), which assesses whether the results in Table 6 are consistent with canonical models of portfolio choice. Explicitly, the regression equation is of the form

$$\Delta Y_i = \beta_0 Middle_i + \beta_1 \left(Middle_i \times High\ Risk\ Tolerance_i \right) + \delta High\ Risk\ Tolerance_i + \tau + u_i,$$

where i indexes household; and *High Risk Tolerance*$_i$ indicates whether i voluntarily chooses a higher risk tolerance score than that recommended by the advisor. The remaining notes are the same as in Table 6.

Table A7: Imputation of Participation in the Stock Market

Imputation Model:	Nonparametric Models		Regression Models	
	Boosted Trees (1)	Random Forest (2)	Logistic (3)	Lasso (4)
New Participants in the Stock Market as a Share of:				
New Middle-Class Robo Participants	0.61	0.57	0.99	0.99
New Upper-Class Robo Participants	0.07	0.01	0.29	0.29
Area under the ROC Curve (Validation Set)	0.99	0.95	0.86	0.86
Area under the ROC Curve (Test Set)	0.96	0.87	0.75	0.75
Accuracy (Test Set)	0.97	0.87	0.77	0.77

Note: This table shows the share of middle-class and upper-class households who previously did not participate in the stock market prior to becoming robo participants after the reduction. Stock market participation status is imputed, and each column corresponds to a separate prediction model. Columns (1)-(2) use nonparametric (i.e., tree-based) models and columns (3)-(4) use models based on logistic regression. Random Forest corresponds to a random forest classification algorithm applied to the SCF dataset, where the set of features are age, income, and liquid assets. Boosted Trees corresponds to a similar boosted trees algorithm. Logistic corresponds to a standard logistic regression. Lasso is a logistic regression with regularizations (i.e., "shrinkage" penalties). The bottom rows of the table report area under the ROC Curve in the validation and test sets as well as accuracy in the test set as an additional performance metric. Appendix C contains further methodological details. Participation in the stock market is defined in Table A3. Middle-class households are defined as belonging to the second or third U.S. quintile of liquid assets ($1k-$42k), and upper-class households are defined as belonging to the fourth or fifth quintiles (>$42k). Liquid assets are defined in Table 1.

Table A8: Imputation of Participation in Asset Management

Imputation Model:	Nonparametric Models		Regression Models	
	Boosted Trees (1)	Random Forest (2)	Logistic (3)	Lasso (4)
New Participants in Asset Management as a Share of:				
New Middle-Class Robo Participants	0.86	0.99	1.00	1.00
New Upper-Class Robo Participants	0.49	0.52	0.65	0.65
Area under the ROC Curve (Validation Set)	0.98	0.92	0.73	0.73
Area under the ROC Curve (Test Set)	0.95	0.82	0.60	0.60
Accuracy (Test Set)	0.96	0.84	0.70	0.70

Note: This table shows the share of middle-class and upper-class households who previously did not participate in professional asset management prior to becoming robo participants after the reduction. Participation in professional asset management is imputed, and each column corresponds to a separate prediction model. Columns (1)-(2) use nonparametric (i.e., tree-based) models and columns (3)-(4) use models based on logistic regression. Participation in professional asset management is defined in Table A3. The remaining notes are the same as in Table A7.

Table A9: Imputation of Homeownership

Imputation Model:	Nonparametric Models		Regression Models	
	Boosted Trees (1)	Random Forest (2)	Logistic (3)	Lasso (4)
Homeowners as a Share of:				
New Middle-Class Robo Participants	0.39	0.44	0.28	0.28
New Upper-Class Robo Participants	0.84	0.93	0.78	0.78
Area under the ROC Curve (Validation Set)	0.98	0.91	0.76	0.76
Area under the ROC Curve (Test Set)	0.95	0.83	0.67	0.67
Accuracy (Test Set)	0.96	0.84	0.71	0.71

Note: This table shows the share of new middle-class and upper-class robo participants who own owner-occupied residential real estate (i.e., homeowners). Homeownership is imputed, and each column corresponds to a separate prediction model. Columns (1)-(2) use nonparametric (i.e., tree-based) models and columns (3)-(4) use models based on logistic regression. The remaining notes are the same as in Table A7.

Table A10: Imputation of Wealth Class

Imputation Models:	Nonparametric Models	
	Boosted Trees (1)	Random Forest (2)
Middle-Class Robo Participants Classified as Middle-Class	0.58	0.61
Upper-Class Robo Participants Classified as Upper-Class	0.62	0.59
Accuracy (Validation Set)	0.90	0.89
Accuracy (Test Set)	0.85	0.84

Note: This table shows the results of wealth class reclassification of new robo participants based on nonparametric (i.e., tree-based) models. Random Forest corresponds to a random forest classification algorithm applied to the SCF dataset, where the set of features are age and income. Boosted Trees corresponds to a similar boosted trees algorithm. For each model and wealth class, we report the share of new robo participants whose wealth class based on self-reported liquid assets equals the wealth class based on imputed liquid assets. The remaining notes are the same as in Table A7.

Trust and Race: Evidence From the Market for Financial Advice

Christopher P. Clifford[a], William C. Gerken[a,*], Tian Qiu[a]

[a]*Department of Finance & Quantitative Methods, University of Kentucky, Lexington, KY 40506*

PRELIMINARY WORKING PAPER, DO NOT CITE WITHOUT PERMISSION

Abstract

In this paper, we examine the relation of racial homophily between financial advisors and local households, and its effect on stock market participation. We find that in ZIP codes that are majority non-white, the presence of a local matched minority financial advisor increases stock market participation. We find this result to be concentrated in middle-income households with more modest effects among high-income households. Our results help explain the persistent racial group differences in stock market participation rates that remain even after controlling for income.

*Corresponding Author. Tel: +1 (859) 257-2774 Fax: +1 (859) 257-9688

Email addresses: `chris.clifford@uky.edu` (Christopher P. Clifford), `will.gerken@uky.edu` (William C. Gerken), `tian.qiu@uky.edu` (Tian Qiu)

August 2, 2021

1. Introduction

A puzzle in the household finance literature is why so many households in the U.S. do not participate in the stock market given the high returns and advantages of a well-diversified portfolio (Mankiw and Zeldes, 1991). One explanation put forth is that participation has both monetary and non-monetary fixed costs (Haliassos and Bertaut, 1995). Indeed, wealth and income levels are positively correlated with participation, suggesting that participation costs may be too high for low income households to justify investment. Yet, even among high income households, there exists significant variation in participation rates that can't be fully explained by fixed costs (Vissing-Jorgensen, 2003).

For example, participation rates among households in the Upper Midwest are substantially higher than those of the Deep South, even within the same income bracket (See Figure 1). Further, there is significant racial variation in participation rates. A study by the Social Security Administration found that even within the same income quartile, White households were twice as likely to invest in financial markets than Black and Latino households.[1] From a policy perspective, such a divide in participation rates can exacerbate gaps in wealth inequality.

We seek to contribute to this literature by studying the role that the racial composition of local financial advisors play in stock market participation. Because financial advisors play a key role in facilitating household access to financial markets, participation is likely driven by the degree to which households can trust their local advisors (Gennaioli, Shleifer, and Vishny, 2015). A 2018 survey by the Certified Financial Planning (CFP) board of hiring managers found that the majority believed that advisors have an advantage with clientele of their own racial and ethnic background. Absent repeated interactions with an individual, investors may form their expectations of trust based on observable signals of social group or background.

[1] A study by Ariel Investments found that 86% of White households with income of at least $50,000 owned stocks or mutual funds, while only 67% of Black households with similar income did.

For example, Stopler and Walter (2019) show that advisor-client homophily affects current clients' willingness to follow their advisors' financial advice. Male clients were more likely to follow the advice if their advisor was of the same gender and age, while female clients were more likely to follow the advice if their advisor had a similar marital and parental status.

In this paper, we ask what role, if any, does racial homophily between client and advisor play in stock market participation? Egan, Matvos, and Seru (2020) report less than 3% of financial advisors are black.[2] Given to the history of wealth distribution in the U.S. and the persisting racial wealth gap, it is not surprising that the financial planning profession currently has a predominantly white clientele base. While 14% of US households are Latino and 13% are black, only 8% and 6% of these households, respectively have incomes over $150k.[3]

Studying this issue presents several empirical challenges. First, it requires access to detailed data for a large sample of financial advisors and the ability to control for other potentially confounding characteristics likely to affect households' decision to participate in the stock market. Second, we need be able to categorize advisors and their potential clients by race. Third, we need to be able to measure investor stock market participation.

To this end, we begin by employing data from the FINRA Brokercheck, Meridian IQ, and Investment Adviser Public Disclosure databases. These three databases allows us to track over 1.3 million unique brokers across over 50,000 firms. The data give dates and branch locations on when individuals are employed with member firms, achieve professional advancement, and engage in professional misconduct. This allows us to not only control for observable characteristics, but to use fixed effects to control for unobservable heterogeneity across time and zip code.

[2]The 2018 survey by the Certified Financial Planning (CFP) board found only 3.5% of CFPs are Black or Latino.

[3]2019 Current Population Survey, US Census Office

Importantly, we can infer advisor race using Bayesian Improved First name Surname Geocoding (BIFSG) of Voicu (2018)[4]. Similarly, we can use detailed Census data (ACS Survey) to measure ZIP code level demographic and socioeconomic characteristics. Finally, we use detailed IRS data at the ZIP code level to examine localized stock market participation. The intersection of these data allows us to measure stock market participation at the ZIP code-year level and study how advisor race matters for trust in financial markets and stock market participation.

2. Preliminary Findings

We find that in ZIP codes that are predominantly non-white (See Figure 2), the presence of an advisor of similar race leads to an increase in stock market participation (See Table 1). For example, in a ZIP code that is predominantly Black, we find that even after controlling for socioeconomic data such as income and education, the presence of a black advisor leads to a 12% increase in stock market participation. To account for unobservable variation that may confound our results, we study how within-zip variation in minority representation affects participation rates. Within the same zip code, we find that the addition of at least one minority advisor leads to a 4% increase in stock market participation for clients of similar race. Our results suggest that, on average, racial homophily plays an important role in facilitating investor participation in equity markets.

Next, we turn to the interaction of race and income on stock market participation. As has been shown in the literature individuals with higher levels of income are more likely to participate in the stock market. In our own data, for example, we find that individuals making \$200K+ (\$50K-200K) are 11× (3×) more likely to participate in the stock market

[4]BIFSG can be viewed as a refinement of the more widely used BISG. While BIFSG has been shown to have modest improvement in classification over Bayesian Improved Surname Geocoding (BISG) in mortgage lending and voting datasets, the largest improvements occur for blacks, a group for which the BISG performance is weakest.

than investors making less than $50,000. Instead, we are interested in how different income levels affect the degree to which race is used a mechanism to form trust. For example, wealth may serve as a proxy for a client's education levels or experience in the financial markets. In this case, the role of racial homophily may be less important in forming trust.

We find that income appears to have a differential effect on how racial homophily and trust are formed (See Table 2). For low income earners, unconditional participation rates are quite low and racial homophily appears to play little role in participation, suggesting that fixed participation costs may be the driving force. For high income earners, however, while participation rates are extremely high, here too we find that racial homophily plays no role in participation. Instead, the effect of racial homophily is strongest among mid-income earners ($50K-200K). Relative to their unconditional rate of participation, minority zip codes are 25% more likely to participate in the market if there is at least one advisor in their area of similar race. For this group of households, the threshold for participation may be low enough to justify participation, yet subjective views of trust may impact their choice. To the extent that racial homophily helps these households to trust more, it may have policy implications on the role of racial diversity in the advisory industry.

The decision of an advisor and client to work together is a two-sided matching game. Firms and their advisor's have strong incentives to seek out wealthy clients, regardless of race, as their own compensation is directly tied to the size of their book of business. Historical wealth differences among groups and homophily preference among mid-income households, could lead to these subgroups underinvesting, missing out on the equity premium, and leading to persistence of group wealth inequity.

References

Egan, Mark, Gregor Matvos, and Amit Seru, 2020, When Harry fired Sally: The double standard in punishing misconduct, *Forthcoming*, Journal of Political Economy.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny, 2015, Money doctors, Journal of Finance 70, 91–114.

Haliassos, Michael, and Carol C. Bertaut, 1995, Why do so few hold stocks?, Economic Journal 105, 1110–1129.

Mankiw, Gregory N., and Stephen P. Zeldes, 1991, The consumption of stockholders and nonstockholders, Journal of Financial Economics 29, 97–112.

Stopler, Oscar, and Andreas Walter, 2019, Birds of a feather: The impact of homophily on the propensity to follow financial advice, The Review of Financial Studies 32, 524–563.

Vissing-Jorgensen, Annette, 2003, Perspectives on behavioral finance: Does "irrationality" disappear with wealth? Evidence from expectations and actions, NBER Macroeconomics Annual 18, 139–194.

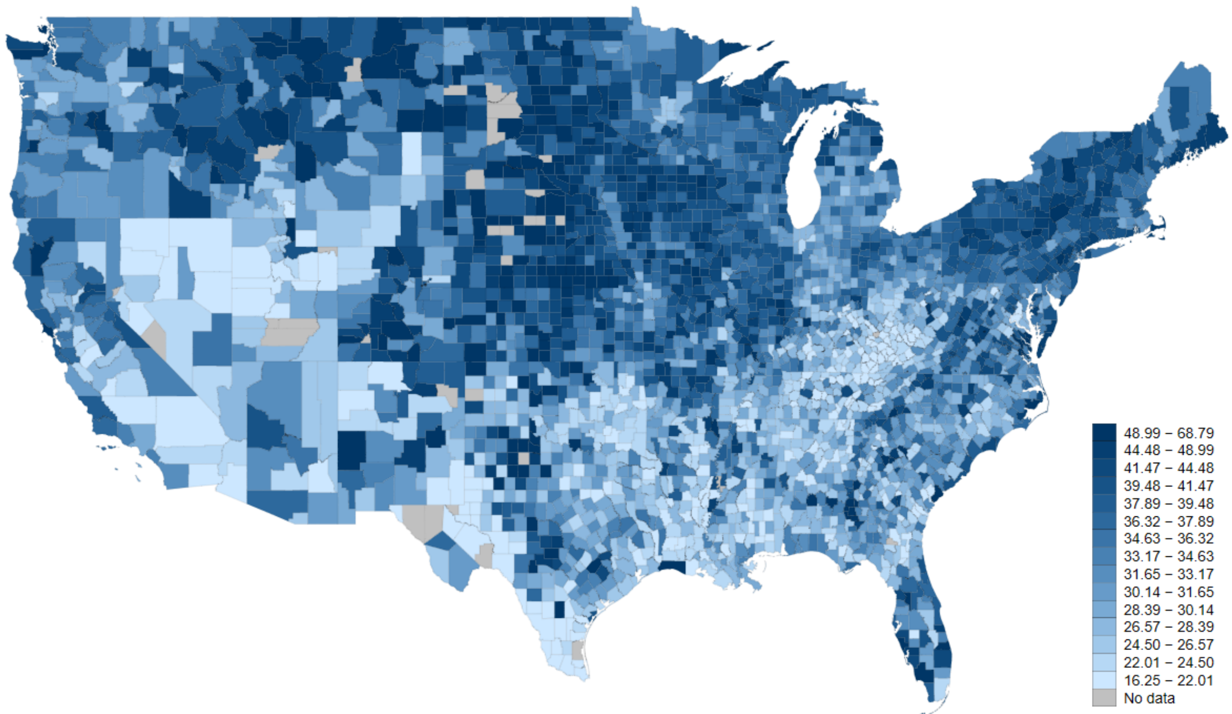

Fig. 1
Equity market participation
This figure documents the geospatial nature of equity market participation for households in the $75k-$100k income bracket. Darker shades of blue represent higher levels of equity participation. Grey represents missing data.

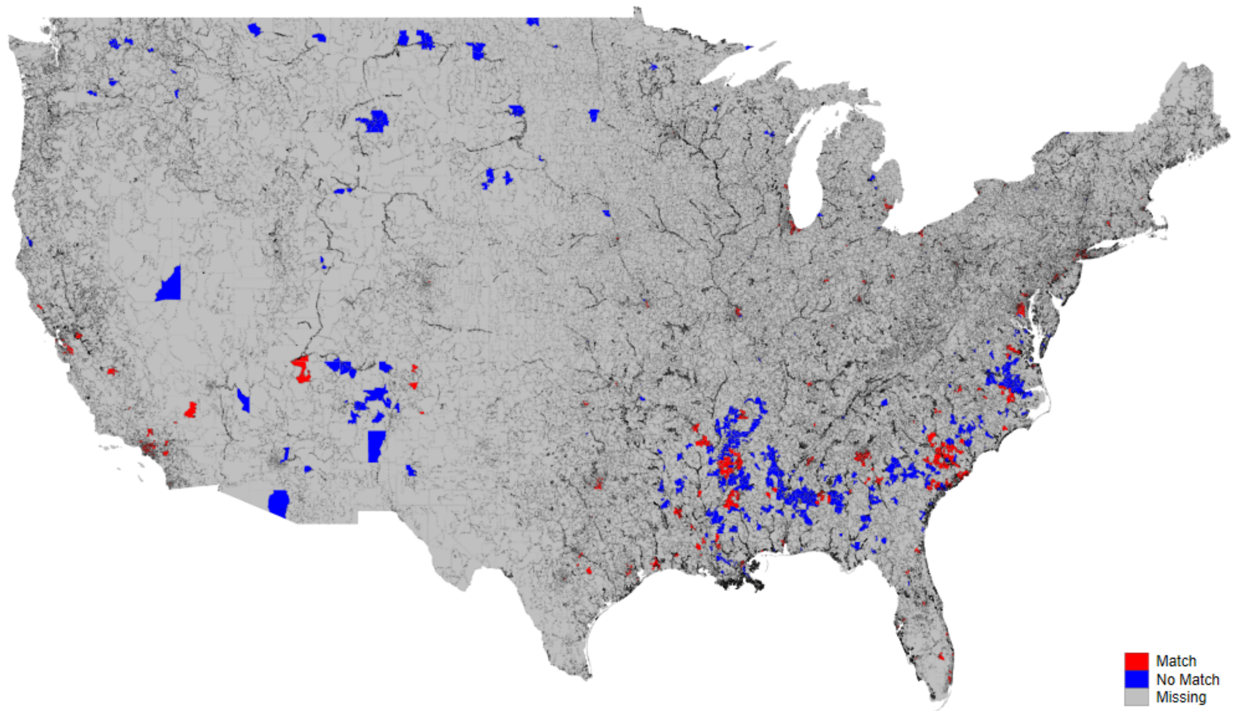

Fig. 2
Distribution of Minority Zip Code
This figure documents the universe of Minority-Majority zips in the US. A zip is Minority-Majority if more than 50% of its total population is represented by a single non-White race group. Red areas are zips with minority financial advisors that match the local majority race group. Blue areas are zips without matching advisors.

Table 1
Racial Homophily and Stock Market Participation

This table presents the effect of matching advisor and local population racial demographic on stock market participation. The unit of observation is a zip code-year. Participation is reported in percentage points. t-stats are reported in the parenthesis. Standard errors are clustered at the zip level.
*** p<0.01, ** p<0.05, * p<0.1

	1	2	3
	Participation	Participation	Participation
Match	1.861***	0.782***	0.276**
	(6.86)	(2.99)	(2.57)
Year FE	Yes	Yes	Yes
Controls		Yes	
Zip FE			Yes

Table 2
Racial Homophily and Income Effects
This table presents the heterogeneous effect of matching advisor and local population racial demographic on stock market participation for households from different income brackets. The unit of observation is a zip code-year. Participation is reported in percentage points. t-stats are reported in the parenthesis. Standard errors are clustered at the zip level.
*** p<0.01, ** p<0.05, * p<0.1

	1	2	3
	Participation	Participation	Participation
50K-200K	9.287***	9.288***	9.167***
	(15.10)	(15.09)	(15.77)
200K+	39.28***	39.31***	32.69***
	(18.01)	(17.90)	(15.48)
Match	0.613***	0.706*	-1.551
	(3.87)	(2.12)	(-1.63)
Match × 50K-200K	2.295***	2.295***	2.400***
	(4.31)	(4.31)	(4.64)
Match × 200K+	-1.403	-1.413	2.765
	(-0.62)	(-0.63)	(1.26)
Year FE	Yes	Yes	Yes
Controls		Yes	
Zip FE			Yes

Where is the Intersection of Madison Avenue and Wall Street? Advertisement, Local Access to Investment Advice, and Stock Market Participation

July 31, 2021

Abstract

We examine the effects of advertisement by investment advisory firms on household stock market participation. Exploiting variation in exposure to financial advertising for households along designated market area (DMA) borders, we find evidence that increased advertising by investment advisory firms leads to higher stock market participation. Importantly, the effects are concentrated in counties where the advertising firm has a physical presence. Consistent with fixed cost frictions to participation, these effects are predominantly among higher income households. We also find larger effects in counties with higher income and racial diversity —areas that are are likely to have lower trust. Our results highlight the complementary nature of persuasive advertising and local access to finance for building trust in the market for investment advice.

Despite the potential benefits of investing in equity markets, household stock market participation rates are low in the U.S. (see e.g., Hung, Clancy, Dominitz, Talley, Berribi, and Suvankulov, 2008; Brown, Ivković, Smith, and Weisbenner, 2008). Households may lack the knowledge and capacity to evaluate investment opportunities (Lusardi and Mitchell, 2007; van Rooij, Lusardi, and Alessie, 2011) or may be too nervous or anxious to make risky investments on their own (Gennaioli, Shleifer, and Vishny, 2015). Instead, professional investment advice may help households overcome these shortcomings and participate in equity markets —if households know about and choose to acquire this advice.

Firms often use advertising to promote awareness and use of their products or services, and firms in the investment advisory industry are no exception. Collectively, financial firms spent $9 billion in advertising in 2019 —an amount that exceeded combined spending in the 2020 presidential election cycle (Kantar Media). Investment advisors typically advertise their services based not on past performance but instead on trust, experience, and dependability (Mullainathan, Schwartstein, and Shleifer, 2008). Thus, it is not surprising that investors cite "trust" as the most important factor when seeking professional advice (Hung, Clancy, Dominitz, Talley, Berribi, and Suvankulov, 2008). Using survey evidence, Burke and Hung (2015) find that trust is an important predictor of the use of professional investment advice. Consistent with the conjecture in Gennaioli, Shleifer, and Vishny (2015) that this trust in investment advisors derives from persuasive advertising, familiarity, and other personal connections, in this paper we examine how this spending on advertisement by investment advisory firms affects stock market participation.

Despite the large sums spent on advertising, understanding the effects of such expenditures is challenging due to the endogenous nature of advertising choices. For example, Jain and Wu (2000) find that mutual funds are more likely to advertise after periods of strong performance, and Mullainathan, Schwartstein, and Shleifer (2008) find mutual funds are more likely to include return metrics in advertisements during bull markets. To overcome such challenges, we take advantage of several unique features of the market for advertising, geographic data on the location of investment firms, and granular geographic data on household stock market participation. A key feature of our approach is that advertising is typically sold at the designated market area (DMA) level. These regional designations were first drawn based on broadcast areas during the birth of television in the 1950s and created arbitrary borders between areas that receive certain advertisements at the same point in time.

Take the example of Suwannee and Columbia counties in northern Florida. Suwannee County was carved out of Columbia County on December 21, 1858. These counties are quite similar in average income, educational attainment, demographics, population density, and climate.[1] Each county is located about a two hour's drive from both the cities of Jacksonville and Tallahassee. However, because of the DMA border running between them, households in neighboring Suwannee county receive one set of advertisements as a result of their placement in the Tallahassee DMA, while households in Columbia county receive a different set of advertisements in the Jacksonville DMA. Advertisement sales in each DMA may be influenced by many factors beyond demand from investment advisory firms including varying demand from local political campaigns, demand from other local industries, and other variation in local economic conditions in the DMA. In our study, we focus on small border counties like Suwanee and Columbia which make up only a tiny fraction of the total households in a DMA, such that these households are unlikely to influence a firm's choice to purchase advertising in a particular DMA. As such, these factors create arbitrary variation in the exposure to investment firm advertisements for households in the border counties.

We exploit this variation in advertising along these market boundaries by implementing a

[1]Suwannee County has an average income of \$23,000, college educational attainment of 16%, average age of 43 years, and a population density of 64 per square mile. Columbia County has an average income of \$26,000, college educational attainment of 15%, average age of 40 years, and a population density of 87 per square mile. Both voted heavily Republican in each of the last six elections.

border discontinuity strategy[2] to identify the effect of advertising on household stock market participation. Using aggregate county-level income tax data from the year 2004 to 2018, we find robust evidence that increased advertising leads to higher stock market participation. In our specifications, we include two sets of fixed effects. $Border \times DMA$ control for fixed characteristics such as demographic differences on each side of a DMA border, though we find little evidence of observable differences given the arbitrary DMA boundaries. $Border \times year$ fixed effects control for any time-varying factors that affect the entire border region on both sides of a DMA border. Across all households, the investment advisory firm advertising elasticities are significant, but somewhat smaller in magnitude than those found in other industries (see e.g., Shapiro, Hitsch, and Tuchman, 2021). However, when we disaggregate households by income tiers, we find that the effects are significantly larger among higher income households: for example, a one standard deviation increase in advertising dollars leads to a 2 percent increase in stock market participation for households with annual income above $200,000. Consistent with models in which fixed costs, such as those related to learning about investment opportunities, can help explain the low participation rates (e.g., Vissing-Jorgensen, 2003), we find little evidence of the effect of advertising among low earners.

We next consider the role that local access to finance has on the efficacy of the advertisement. After a household watches an advertisement and makes a decision about whether there is interest in an investment service, the potential client must take further action to obtain the service. Investment advice is inherently a local business. Geographic proximity between client and advisor can enhance the familiarity and other personal connections that Gennaioli, Shleifer, and Vishny (2015) argue are important for establishing trust. Not surprisingly, both major industry regulatory websites, FINRA Brokercheck's and the SEC's Investment Adviser Public Disclosure, allow potential clients to search for advisors within

[2] We follow a methodology similar to that used in recent papers that examine the effects of advertising on the demand for e-cigarettes (Tuchman, 2019), antidepressant drugs (Shapiro, 2018), and consumer packaged goods (Shapiro, Hitsch, and Tuchman, 2021).

a given city or ZIP code. Similarly, the CFP Board, an industry standard organization, defaults to geographic focused searches on its website.

Our data allow us to exploit variation in the geographic presence of investment advisory firms within the border county sample. Specifically, some counties have a branch office of an investment firm that engages in local advertising during a year, other counties contain branch offices of investment firms that choose not to advertise in that DMA during the year, and yet other counties do not have a branch office of an investment firm at all. This allows us to separately estimate the joint effect of having a local branch of a firm that advertises in a given year from the effect of simply having a local investment firm physically present in the county.[3]

Again, an example may help illustrate the approach. Hamilton County borders Suwanee County and likewise lies within the Tallahassee DMA. Baker County borders Columbia County and likewise lies within the Jacksonville DMA. Fidelity has branch offices in Suwanee and Baker counties, but not in Columbia and Hamilton. Thus, we are able to observe both counties with a local firm presence and those without. For those counties that have a local firm presence, whether the local firm advertises depends on Fidelity's decision to advertise in the Tallahassee and Jacksonville DMAs. Importantly, the Baker County Fidelity branch is only one of 40 Fidelity branches in the Jacksonville DMA, while the Suwannee County Fidelity branch is only one of 10 in the Tallahassee DMA. Thus it is unlikely that either of these remote branch locations are likely to the motivating factor for advertising in their respective DMAs.

We document that the effects of advertising are concentrated in those counties in which advertising firms have a local presence. The magnitude of the effect of advertising on participation in counties where a local firm advertises is nearly double in magnitude among

[3]Our approach assumes that potential clients who are exposed to advertising do not cross the DMA border and purchase in the adjoining DMA. Violations of this assumption will bias the estimates towards zero.

high income households. For comparison, this effect size is similar in magnitude to the effect of community effects documented in Brown, Ivković, Smith, and Weisbenner (2008). These results highlight the multifaceted nature of trust in the investment advisor industry as persuasive advertising is most effective in encouraging participation when complemented by local access to financial professionals.

If persuasive advertising helps to establish trust between local advisors and clients, then one might expect larger effects in areas with lower levels of trust. Using data from the General Social Survey, Alesina and La Ferrara (2002) document lower levels of trust among those living in a racially mixed community or in communities with a high degree of income disparity. Consistent with the findings of Guiso, Sapienza, and Zingales (2008), we indeed find lower participation in such areas; yet, we document that in these areas advertising is more effective in encouraging stock market participation relative to areas of high trust. These findings are consistent with advertising overcoming traditional barriers to participation through developing trust.

While the research design rules out most obvious confounding factors, one remaining concern is that that parallel-trends may not hold, invalidating the difference-in-differences design. We check observable characteristics including number of investment advisors on each side of the DMA borders and see no significant differences. We perform additional placebo tests. Using our main specification, we replace the dependent variable with average salary among households, helping to rule out the possibility that the association was spuriously caused by some time-varying local economic shocks. We then use our main specification and replace the independent variable with other non-investment financial advertising (advertising by and for credit unions, loans, and lending) to help rule out other spurious time-varying advertising shocks.

While the border sample aids the identification of the effect, the counties may not be representative and raise concerns about generalizability to the larger population. To address this, we run a naïve county-year panel regression using data from all counties in the US (i.e. across all 210 DMAs). While this approach sacrifices careful identification, we demonstrate that the method produces a similar estimate suggesting that the relationship between advertising and participation should generalize.

Our study contributes to the literature on the determinants of stock market participation. It is well documented that stock market participation is low (Brown, Ivković, Smith, and Weisbenner, 2008; Haliassos and Bertaut, 1995; Hong, Kubik, and Stein, 2004), and that a lack of trust (Burke and Hung, 2015; Guiso, Sapienza, and Zingales, 2008) and insufficient financial knowledge (Hilgert, Hogarth, and Beverly, 2003; Lusardi and Mitchell, 2007) are contributing factors to these low rates. It also relates to the growing literature about how the access to finance affects consumer financial well-being. Brown, Cookson, and Heimer (2019) find that local provision of finance matters for consumer financial credit outcomes. We present causal evidence that advertising investment advice is an effective means of addressing these shortfalls and spurring individual action to participate in markets. To our knowledge, this paper is the first to directly explore this question.

We also contribute to the broader literature on the efficacy of advertising in the financial services industry in general. Much of the literature on advertising in this space focuses on mutual funds despite the challenge that the endogenous advertising choice presents in establishing causality between advertising and flows. Jain and Wu (2000) examine a small sample of funds and find that after controlling for past performance, funds that advertise have higher flows than a group of funds matched on investment style and prior performance. Similarly, Gallaher, Kaniel, and Starks (2015) demonstrate that advertising weakens the sensitivity of the flow-performance relationship for poorly-performing funds while advertising increases the flow-performance relationship for top-performing funds. Barber, Odean, and Zheng (2005) find evidence that while mutual funds fees are associated with outflows, the negative effects of fees on flows is offset when the fees are spent on marketing. We supplement these studies by similarly documenting efficacy of advertising in the investment advisory space but employ an approach that overcomes the many issues of endogenous selection inherent to the current literature.

This paper is also related to literature about the persuasive view of advertising in financial products and services (e.g., see Hortaçsu and Syverson, 2004, and Elton, Gruber, and Busse, 2011, for the mutual fund industry; Gurun, Matvos, and Seru, 2016; Allen, Clark, and Houde, 2014, 2019, for mortgages; Green, Hollifield, and Schürhoff, 2007, for bonds; Duarte and Hastings, 2012, for privatized social security plans; Christoffersen and Musto, 2002, for money funds; and Brown and Goolsbee, 2002, for life insurance). The literature typically focuses on persuasion that creates "artificial" product differentiation. We differ in that we find evidence of the persuasive view of advertising that creates firm-specific ties that enable investors to take risk, consistent with the model of Gennaioli, Shleifer, and Vishny (2015). While costly, these expenditures help to build trust that enables households to participate in markets.

This paper proceeds as follows: in Section 1, we describe our identification strategy in detail. In Section 2, we summarize the advisor, advertising, and income data sets and define the construction of our main dependent variable. Section 3 presents our main results of the effects of advisor advertising on stock market participation, and includes placebo tests. In Section 4, we consider to what extent the interaction between advertising and the local presence of a financial advisor has on participation, and explore the varying effects of advertising on stock market participation due to county heterogeneity. Finally, Section 5 concludes.

1 Institutional Background and Research Strategy

Identifying a causal relationship between investment advisor advertising and stock market participation can be challenging. Advertising is not randomly assigned: financial advisory firms may target ads where residents are more likely to be stock market participants. Ordinary least squares estimates of the effects of advertising on action will lead to a positive bias if we fail to account for firms choosing to advertise where residents are more likely to trade. To overcome this endogeneity problem, we employ the methodology of Shapiro (2018) and Tuchman (2019) and take advantage of geographic discontinuities in local advertising markets. Specifically, we examine how demographically similar households on opposite sides of advertising market borders respond to financial advisor advertising.[4]

The Nielsen Company (Nielsen) assigns groups of counties in exclusive geographic areas to designated market areas (DMAs). The DMA to which a county is assigned determines what local programming a household within the DMA's borders will receive. FCC regulations and federal law in most cases prohibit broadcasters, including cable and satellite operators, from providing signal to a home in a DMA other than the signal to which homes in that DMA are meant to receive. Therefore, local television ads are transmitted to all households within a particular DMA, and ads purchased for a particular DMA only may not be shown in other DMAs.[5] Similarly, DMA-level advertising takes places at the magazine, newspaper, radio, and outdoor levels.

A useful demonstration is that of television commercials during the Super Bowl. Over 100 million viewers tune in to the NFL's annual championship game which commands the largest U.S. viewing audience each year. Naturally, ads aired nationally during the game garner much attention, and spots have cost over $5 million for 30 seconds. As they typically

[4]This section draws heavily from the descriptions of the identification strategy in Shapiro (2018) and Tuchman (2019).

[5]See Shapiro (2018) and Spenkuch and Toniatti (2018) for a thorough discussion of designated market areas.

do, networks also allot time during the broadcast to local ads.[6] In these slots, advertisers purchase slots to target ads to specific regional campaigns. For example, the Church of Scientology purchased time in 16 DMAs, including those where it has major centers. In the Lexington DMA, the local Winstar Stud Farm touted the prowess of its stallions in a concurrent commercial. Importantly, households in different regions watching identical programming at the same time received different advertisements.

Figure 1 presents a map of the 210 DMAs across the United States. DMAs are typically centered around metropolitan areas and include an average of 15 counties. Counties are assigned to only one DMA and are rarely reassigned. Nielsen originally drew boundaries to reflect demand for television programming while accounting for broadcast television airwave range. As such, boundaries do not reflect political jurisdictions: nearly half of DMAs cross state borders, with an average of 1.7 states represented per DMA. This in part alleviates the concern that the relationship between local advertising and stock market participation we study will be biased by unobserved variables pertaining to local laws and individuals choosing to live on one side of a border.

We obtain identification by comparing advertising and its effects on stock market participation for counties on opposite sides of a DMA boundary. The map of the states of Florida and Georgia presented in Figure 2 demonstrates this approach. The Tallahassee DMA consists of 10 counties in Georgia and 9 counties in Florida.[7] The Jacksonville DMA consists of 6 counties in Georgia and 9 counties in Florida. A total of 7 counties lie on the border of these two DMAs, with Clinch, Echols, Hamilton, and Suwanee counties on the western side of the border, and Ware, Baker, and Columbia counties on the east. Under the assumption that, for example, residents of Suwanee County and Columbia County are

[6] "Local Advertising Was Also On Display During Super Bowl XLIX". Kantar Media Report, 2015.

[7] Depending on the state, local administrative jurisdictions may be referred to as parishes, counties, boroughs, and/or independent cities. We use the terms "counties" to refer to all of these functionally similar regions.

similar on unobservables but exposed to different levels of local financial advisor advertising due to placement in different media markets, we can use their location on a DMA border as a setting for a natural experiment. A necessary condition of this framework is that there is limited selection by individuals more likely to be stock market participants to reside on certain sides of DMA borders. We test this possibility and discuss the results in Section 3.1.

This approach relies on variation in advertising both *across* DMAs and over time *within* a DMA. We graphically depict these distinct sources of variation for the Boston and New York City markets in Figure 3. Each point marks, for a given year, the units and average rate per unit of advertising (in dollars). Panel A shows the downward sloping demand curves in each of these markets: as on might expect, demand for units of advertising falls as the spot advertising rate increases. As apparent from Panel A, there is variation in this relationship within the Boston market and within the New York City markets. Likewise, there is significant variation across the Boston and New York due to distinct time-invariant characteristics inherent to these markets. Panel B of Figure 3 shows the demand curve in these markets with the inclusion of market fixed effects. Again, the demand for advertising in total units falls as the spot advertising rate increases, but here we control for those time-invariant market-specific characteristics. The variation in advertising is thus driven by the competition among advertisers for limited advertising spots.

Given these sources of variation, we incorporate two sets of fixed effects in our main specifications: *border-DMA* and *border-year* fixed effects. A "border" includes all adjacent counties on both sides of two neighboring DMAs' boundary. *Border-DMA* fixed effects permit time-varying advertising intensity *within* the border counties on one side of the border. For example, this allows systematically different levels of demand for advertising for Clinch, Echols, Hamilton, and Suwannee counties of the Tallahassee DMA relative to Ware, Baker, and Columbia counties of the Jacksonville DMA. This strategy relies on the assumption that stock market participation in counties across a DMA border would follow parallel trends in the absence of advertising exposure variation. *Border-year* fixed effects restrict this assumption of common trends in advertising to the local counties of a DMA border, in this case the seven counties Clinch, Echols, Hamilton, Suwannee, Ware, Baker, and Columbia on both sides of the Tallahassee-Jacksonville DMA border.

Firms likely target advertising to the populous metropolitan areas often at the center of DMAs rather than the relatively rural counties on DMA borders. This suggests that the different levels of advertising on DMA borders is less likely driven by characteristics of households in the border counties themselves. Continuing the example from Figure 2, both the Tallahassee and Jacksonville DMAs are named after the major metropolitan areas within their borders. The city of Tallahassee lies off the Tallahassee-Jacksonville DMA border in Leon County, which alone represents 39.7% of the DMA's 2019 population. Likewise, Jacksonville lies in Duval County, which represents 48.5% of the Jacksonville DMA's 2019 population. In contrast, the counties on the Tallahassee-Jacksonville DMA border represent a combined 8.8% and 6.9% of their respective DMA's population in 2019. Following the framework of Li, Hartmann, and Amano (2020), we limit our sample to these "small borders" where the combined population of all counties on each side of the DMA border is less than 10% of the total population in its respective DMA. On these borders, it is more plausible that local demand shocks in these relatively sparse counties are less likely to influence the advertising choices of financial advisors at the DMA level. Our final sample consists of 381 such borders (consisting of 1,155 unique counties), with each a setting for a separate experiment over our sample period.

2 Data

2.1 Registered Investment Advisors

The Investment Advisers Act of 1940 (IAA) defines an investment adviser as any person or company that is compensated for advising others on investing in securities. The IAA authorizes the Securities and Exchange Commission (SEC) to regulate and require registration by investment advisors. As part of their registration, registered investment advisors (RIAs) must file at least once annually a Form ADV with the SEC. Generally, only advisors with at least $100 million in assets under management and 15 or more U.S. clients must satisfy this registration and record-keeping requirement with the SEC.

We collect data on RIAs from their Form ADV, available on the SEC's *Investment Adviser Public Disclosure* website.[8] Because advisors may file more than one ADV per year as a result of material changes, we use the ADV filing active at the beginning of each calendar year from 2004 to 2018. We identify advertising firms based on their presence in the Kantar Media database, which we describe in the next subsection. No unique identifier links the Form ADV investment advisor data to the Kantar advertising data, so we manually match advisors to Kantar based on name. Of the 26,909 financial advisors that file a form ADV at any point over the period 2004-2018, approximately 7.6% or 2,039 RIAs are advertisers and present in the Kantar Media database.

Table 1 summarizes this sample. As one might expect, advertising RIA firms are significantly larger than non-advertising firms, with an average of $35.4 million more in assets under management (t-stat = 5.38). Advertising firms are similarly larger both in their number of accounts and number of employees. RIAs that advertise are also more likely to be retail focused: 54.8% of advertising firms mostly serve individuals and high net worth indi-

[8]Bulk downloads of historical ADVs are available at https://www.sec.gov/foia/docs/form-adv-archive-data.htm.

viduals, relative to the 48.5% of non-advertisers that primarily serve these groups. Finally, advertisers are significantly more likely to offer financial planning services.

2.2 Advertising

Advertising data come from Kantar Media (Kantar). The data include monthly advertising expenditures and advertising units for financial advertisers from 2004 to 2018 across television, print, radio, digital, and out-of-home (billboards, taxis, malls, cinema) channels. Spending and unit data is presented at the DMA and national levels.

Table 2 presents summary data on the Kantar dataset. The unit of observation is a border-county-year. *Total Ad Units* is the number of local television, print, radio, digital, and out-of-home advertisements in small border counties, on average 31,008.3 from 2004 to 2018. *Total Ad Dollars*, the total amount of RIA local advertising spending in a small border county in a year, averages $7.7 million, but is substantially skewed (the median is $1.45 million). The RIA advertising spending per household (*Ad $ per household*) is defined as an RIAs' national advertising expenditure scaled by the national population plus their local advertising expenditure scaled by the population of the DMA. In a small border county in a year, *Ad $ per household* averages $26.53, while the average number of RIAs in a small border county (*# of RIAs*) is 6.4.

2.3 Stock Market Participation

We use data from the IRS's Statistics of Income (SOI) and follow Brown, Ivković, Smith, and Weisbenner (2008) to develop our proxy for stock market participation. SOI provides annual data on tax returns in a zip code, including the number of filers, average reported salaries and wages, average adjusted gross income (AGI), number of returns filed, and number of returns reporting dividends. Following Brown, Ivković, Smith, and Weisbenner (2008), we proxy for stock market participation based on whether an individual reported dividend income on their federal tax return in a given year. That is, for county i in year t:

$$Participation_{i,t} = \frac{Number\ households\ reporting\ dividend\ income_{i,t}}{Tax\ returns\ filed_{i,t}} \tag{1}$$

Brown, Ivković, Smith, and Weisbenner (2008) note two weaknesses of this measure: (1) dividend income reported on tax returns includes distributions from any mutual fund, including those solely invested in fixed income and (2) it may not capture those who invest exclusively in non-dividend paying firms, such as growth-oriented stocks. To the extent that we wish to measure individuals' participation in financial markets overall, (1) does not present a serious issue, and (2) is not an issue so long as an investor holds any *one* dividend paying stock or mutual fund that makes a distribution. Brown, Ivković, Smith, and Weisbenner (2008) additionally show that the correlation between actual equity market participation as reported by the FED's Survey of Consumer Finances across four pooled cross-sections and their measure was 0.62. Lin (2020) finds that this measure of participation and the measure of participation from the University of Michigan's Health and Retirement Study used in Hong, Kubik, and Stein (2004) is 0.69.

Table 3 presents summary statistics on the SOI data by AGI in 2018. In Panel A, we present the full United States sample. There were 148.2 million tax returns filed in 2018. Our stock market participation measure is 24.8% overall, and monotonically increases from an average of 10.9% in the $1,000 - 25,000 AGI tier to 56.6% for households with AGI above $200,000. As we will describe in detail in the next section, we follow the framework of Li, Hartmann, and Amano (2020) and limit our sample to counties on "small DMA borders," or borders where the combined population of all counties on each side of the DMA border is less than 10% of the total population in its respective DMA. Panel B of Table 3 summarizes the households in these counties, with similar rates of participation across AGI tiers as in the full United States sample.

3 Advertising and Stock Market Participation

3.1 Baseline Tests

In our baseline tests, we regress our measure of stock market participation on advertising in the following fixed effects specification:

$$Participation_{i,t} = \beta Advertising_{d,t} + \delta_{b,d} + \delta_{b,t} + \epsilon_{i,t} \tag{2}$$

Subscripts are defined as: county i, year t, DMA d, and DMA-border b. *Participation* is the fraction of total household tax returns that report dividend income at the county-year level. We measure *Advertising* in three ways: (1) the natural logarithm of the units of RIA local advertising in a DMA-year, (2) the natural logarithm of dollars of RIA advertising expenditures in a DMA-year, or (3) *$ per household*, or RIAs' national advertising expenditure scaled by the national population plus their local advertising expenditure scaled by the population of the DMA following the approach in Shapiro (2018). $\delta_{b,d}$ represents the border-DMA fixed effects while $\delta_{b,t}$ represents the border-year fixed effects.

Table 4 presents the results of our baseline specification. Following the framework of Abadie, Athey, Imbens, and Wooldridge (2017), we cluster standard errors by DMA, the level of the assignment mechanism. All columns report coefficient estimates with the inclusion of the border-DMA and the border-year fixed effects. In columns (1) and (2), we find a positive and statistically significant relationship between stock market participation and local advertising both in financial advisor advertising units and financial advisor advertising expenditures. Given our specification and the inclusion of the border-DMA and border-year fixed effects, column (1) implies stock market participation in a county increases as a result of greater advertising intensity relative to local advertising on that county's side of the border ($\delta_{b,d}$) and relative to local advertising across all counties on both sides of the DMA border in that year ($\delta_{b,t}$).

In column (3) of Table 4, we take as our relevant measure of advertising the *scaled* dollars, or financial advisors' national advertising expenditure scaled by the national population plus their local advertising expenditure scaled by the population of the DMA. We include both a linear and squared term to highlight the diminishing marginal effects of advertising on participation. The scaled dollars measure of advertising is positive and significant while its squared term is negative and significant. Thus, the marginal effects of advertising on stock market participation is lower as the dollars of advertising per household gets increasingly large.

The coefficient estimates from column (1) imply that a one standard deviation increase from the average in local financial advisor units of advertising leads to an increase in stock market participation of 0.30% ($= 0.0016 \times \ln(49171.8 \div 31008.3) \div 0.25$) relative to the average rate of stock market participation. Likewise, for the dollars of advertising, a one standard deviation increase from the average in the dollars spent on DMA advertising leads to an increase in stock market participation of 0.39% ($= 0.0011 \times \ln(18900000 \div 7734262) \div 0.25$). This average advertising elasticity across all households is smaller in economic magnitude than the ones for consumer packaged goods (see e.g., Shapiro, Hitsch, and Tuchman, 2021).

We additionally test whether there are systematic differences in characteristics of county residents lying on either side of a DMA border. If, for example, individuals with higher income were more likely to cluster on the side of a DMA border with greater levels of advertising, it may bias our results if those individuals are more likely to be stock market participants. To test this, we compute the average number of financial advisors (scaled by the population), income, population, education attainment, debt-to-income ratios, social capital index, and age on both the sides of a DMA border. The number of advisors comes from FINRA BrokerCheck data. Income data is from IRS SOI. Population and educational attainment is from Census.gov. Debt-to-income ratios are from the Federal Reserve (federalreserve.gov), and the Social Capital Index is from the United States Congress Joint Economic Committee "The Geography of Social Capital in America" (jec.senate.gov). We consider the side of the border with greater (less) advertising to be the *Heavy* (*Light*) side of our small borders sample. We average across counties and aggregate to the border-DMA level.

Table A1 presents these results. On average, there are 10.06 financial advisors per 10,000 people on the side of the border with greater advertising and 10.55 financial advisors per 10,000 people on the side of the border with relatively lower levels of advertising. This difference is not statistically significant (t-stat = -0.50, p-value = 0.616). Similarly, the difference in average income, population, percent of residents with a college degree, debt-to-income ratios, social capital index, and age of county residents on the heavy and light sides of DMA borders cannot be shown to be statistically different from zero. These similarities across borders helps to alleviate the concern that individuals who are more likely to be stock market participants select onto the side of a DMA border with greater levels of advertising.

3.2 Adjusted Gross Income

Recalling Table 3, stock market participation varies substantially across income ranges. Overall, a quarter of households participate in trading, but only about 1 in 10 households with income less than \$25,000 invest while more than half of households with greater than \$200,000 in income invest. We next explore how advertising effects stock market participation among households with varying levels of income.

In order to determine the marginal effects of advertising on participation across households of disparate income, we interact our measure of advertising with income ranges, denoted with subscript k in the following specification:

$$Participation_{i,k,t} = \sum_{k} \beta_k (AGI\ Tier_k \times Advertising_{d,t}) + \delta_{b,d} + \delta_{b,t} + \delta_{k,t} + \epsilon_{i,t} \qquad (3)$$

As in our baseline specification presented in Equation 2, we include border-DMA $\delta_{b,d}$ and border-year $\delta_{b,t}$ fixed effects. Additionally, we incorporate AGI tier-year fixed effects $\delta_{k,t}$ to achieve identification of advertising on participation *within* an AGI tier and control for heterogeneity across households in different income groups.

Table 5 presents the results of Equation 3 and demonstrates that a positive and statistically significant relationship between advertising and participation is present for higher income levels only. Column (1) interacts the AGI tier with the natural logarithm of financial advisor units of advertising in the county's DMA, while column (2) interacts the AGI tier with the natural logarithm of the dollars of financial advisor advertising expenditures in the county's DMA. Among wealthier households, we find that the effects are significantly larger than our baseline effects. But we see no significant relationship between advertising and participation in either units or dollars of advertising for lower income households. In part, this contextualizes the economic magnitude of our estimates in Table 4: the relationship is driven by households that combined represent only approximately 20% of the total tax-filing households in the United States.

3.3 Retail-Focused RIAs

RIAs vary substantially in the services they provide. In our sample, about half of the RIAs primarily advise mutual funds, hedge funds, banks, or endowments while about half primarily offer planning and retirement services to individual investors. We hypothesize that advertising by RIAs with a retail focus will have effects on stock market participation. We test this in Table A2. We identify *Retail RIAs* as any RIA where at least 50% of its clients are either individuals or high-net worth individuals as reported on their Form ADV. From column (1), the coefficient on *Log(Units of Retail RIA Advertising)* is statistically significant at the 5% level. Thus, stock market participation increases as the units of advertising for retail-oriented RIAs increases. This relationship is not present in RIAs without a retail focus, as the coefficient on *Log(Units of Non-retail RIA Advertising* is not significant. Additionally, we reject that the difference between these two effects is zero.

Similarly, column (2) of Table A2 shows that the dollars of advertising by retail-oriented investors increases stock market participation while dollars spent on advertising by non-retail RIAs does not. The difference between these coefficients is different from zero. Taken together columns (1) and (2) of Table A2 support a role for trust in advertising and relationship between individuals and their investment adviser.

3.4 Alternative Specifications

In this section, we show the robustness of our baseline results with alternative advertising measures, sample selection, and border implementation. First, our identification strategy relies on the discontinuity of advertising across DMA borders. This identifying assumption could be violated if not all types of advertising are bought at the DMA level. To address this concern, we follow Shapiro (2018) and Tuchman (2019) and repeat our main analysis using television advertising spending only. This result is presented in Table A3, column (1). We find results similar to our main analysis using all adverting types.

Second, we determine whether our results hold in a naive panel setting where we include only county and year fixed effects. In this setting, we do not restrict the sample to border counties, nor control for border-DMA or border-year fixed effects. This specification may lack the clean identification of the border strategy, but it offers broader applicability by including all counties. The results presented in column (2) of Table A3 are similar to our main specification.

Finally, we explore an alternative strategy where we pair counties on a DMA border. As opposed to our main tests where we aggregate all border counties on one side of a border, this alternative implementation allows us to include more small counties in the analysis. Rather than limit our analysis to those counties on a border that combined have a population share less than 10% of their DMA, we restrict our sample to individual county pairs that are each less than 3% of their respective DMA population. We additionally include county-pair and pair-year fixed effects. In column (4) of Table A3, we show that this estimate is similar to what we find in Table 4.

3.5 Placebo Tests

Our identification strategy relies on the assumption that stock market participation in counties across a DMA border would follow a parallel trend in the absence of differences in advertising exposure. To mitigate the concern that this parallel trend assumption is violated, we conduct two sets of placebo tests.

First, we replace the dependent variable $Participation_{i,t}$ in our baseline specification presented in Equation 2 with the average salary in a county-year. If there are no differences in economic conditions across DMA borders, we should observe no advertising effects on average salary. Columns (1) and (2) of Table 6 present these results. Neither the units nor dollars of RIA advertising in a county-year have a statistically significant effect on the average salary of residents of a county.

Second, we want to be sure that our main effect is not driven by differences in general advertising across DMA borders. To do so, we replace the main independent variable, RIA advertising $Advertising_{d,t}$, in our baseline specification presented in Equation 2 with measures of other forms of financial firm advertising in a DMA-year. If our main effect is driven by different advertising trends across borders, we should see effects for these types of advertising on stock market participation. Table 6 presents these results. Column (3), (4), and (5) respectively show that the units of credit union advertising, loan advertising, and lending advertising by other financial firms have no statistically significant effect on stock market participation.

4 Advertising and Access to Local Investment Advice

We next consider the role that local access to finance has on the efficacy of the advertisement. After a household watches an advertisement and makes a decision about whether there is interest in an investment service, the potential client must take further action to obtain the service. Investment advice is inherently a local business. Geographic proximity between client and advisor can enhance the familiarity and other personal connections that Gennaioli, Shleifer, and Vishny (2015) argue are important for establishing trust.

In this section, we test whether the *local* presence of a RIA advertiser encourages stock market participation by *local* households. Additionally, we determine to what extent county heterogeneity results in disparate effects in the response by locals to the advertisement of financial advice.

4.1 Local Investment Advisors

We obtain detailed advisor location data from FINRA and the SEC. We then identify the local branches of firms that choose to advertise in a particular DMA, and also those that do not. To examine the extent to which the local access to investment advice has on stock market participation, we regress our proxy for stock market participation in county i in year t on an indicator for the presence of any local RIA and an indicator for the presence of any local advertising RIA:

$$Participation_{i,t} = \beta_1 AnyLocalRIA_{i,t} + \beta_2 AnyLocalAdvertisingRIA_{i,t} + \delta_{b,d} + \delta_{b,t} + \epsilon_{i,t} \quad (4)$$

As in our baseline specification presented in Equation 2, we include border-DMA and border-year fixed effects. These results are presented in Panel A of column (1) in Table 7. The coefficient on *Any Local RIA* is not statistically significant: the physical presence of one or more RIA branches in a county alone does not appear to encourage residents of that county to participate in the stock market. Yet, the coefficient on *Any Local Advertising RIA* is positive and statistically significant, indicating that the presence of any local financial advisory firm that has advertised in that county's DMA in that year does indeed have an incremental and positive effect on stock market participation in that county.

In column (2) of Panel A, Table 7, we substitute the binary variable for any local advertising RIA's presence in column (1) with continuous counts of the number of advertising units of local and non-local RIA advertising in a county. There is a positive and statistically significant coefficient on *Log(Units of Local RIA Advertising)* but no loading on *Log(Units of Non-local RIA Advertising)*. We further confirm that the difference between the coefficients on the *Log(Units)* if local and non-local RIA advertising is different than zero. Similarly, when we examine the dollars of advertising spent by local and non-local RIAs in column (3), we find increased spending by RIA's with a local presence has an effect on participation while spending by non-local RIAs does not. Again, we find the difference between these two coefficient estimates is statistically significant. The collective results presented in Table 7 are consistent with trust building in the advertising of financial advice.

Given the the relationship between advertising and stock market participation is concentrated in high income households (Table 5), we next consider the interaction between local units and dollar spending of advertising with AGI tiers in Panel B of Table 7. In column

(1), we interact AGI tiers with local RIA advertising units and show positive and significant effects for households with annual income in excess of $75,000. There is no statistically significant relationship between local advertising and participation among lower household income tiers. In column (2), we show this relationship similarly holds for dollar spending on local advertising.

In sum, we find compelling evidence that the effects of advertising are concentrated in those counties in which advertising firms have a local presence. Moreover, we find muted effects of advertising in areas without a local presence. These results highlight the multifaceted nature of trust in the investment advisor industry as persuasive advertising is more effective encouraging participation when complemented by local access to financial professionals.

4.2 Community Trust and Participation

Next, we consider heterogeneity in trust across communities. Guiso, Sapienza, and Zingales (2008) present evidence that areas with higher levels of trust are more likely to participate in the stock market. Alesina and La Ferrara (2002) find that high income and racial disparities are among the strongest determinants of low trust. We hypothesize that in areas with low trust, advertising will have a greater effect in encouraging stock market participation insofar as advertising investment advice plays a role in establishing trust between potential clients and financial advisors.

To test this, we use Census data to compute income and racial disparity at the county level by the normalized Herfindahl index.[9] We classify a county as having *High Income Disparity* if the county's normalized Herfindahl index for AGI tiers is above the median for all counties in a given year. We classify a county as *High Racial Disparity* if the county's

[9]We compute the disparity indices as $H = \sum_{i=1}^{N} s_i^2$, and then normalize as $H^* = \frac{H-1/N}{1-1/N}$ for $N > 1$ and $H^* = 1$ for $N = 1$. For income disparity, s is the ratio of tax filers in an AGI tier to the total number of tax filers in that county-year, and N is the number of AGI tiers in that county-year. For racial disparity, s is the ratio of a race group to the total population in the county year, and N is the number of race groups in that county year.

normalized Herfindahl index for race groups in a given year is above the median for all counties in that year. Crucially, we control for the ratio of each race group in each county, to pick up the homogeneity in race, not the variation due to composition for *specific* races. We take as our dependent variable stock market participation for households with AGI in excess of $100,000 given the results in Panel B of Table 7 demonstrating the effects of advertising on participation are concentrated in households with high income.

Table 8 presents these results. As expected areas with low trust as proxied by *High Income Disparity* or *High Racial Disparity* have lower levels of participation consistent prior studies such as Guiso, Sapienza, and Zingales (2008). Households from communities with low trust are less likely to participate in markets. To examine our assertion about the interaction between low trust and advertising, we include interactions between our low trust proxies and local RIA and local advertising RIA indicators. In column (1), we show the interaction between *High Income Disparity* and the presence of a local advertising RIA is positive and statistically significant. In column (2) we show a similar relationship between stock market participation and the interaction of *High Racial Disparity* and local advertising by RIAs. While households in counties with high racial disparity are less likely to participate in the stock market, in these are of low trust, advertising is more likely to encourage participation.

5 Conclusion

In this paper, we demonstrate that advertising by financial advisers increases stock market participation. This relationship is stronger for high-income households and in counties that have a local advertising advisor present. Strength of the community matters for stock market participation, and we show that advertising may help to establish trust in areas with weaker community ties. We achieve identification in the relationship between advertising and participation by comparing differences in participation rates across borders of otherwise similar counties in different designated market areas. Our incorporation of a rich set of fixed effects at the border-DMA and border-year level exploit variation in advertising both *across* and *within* regions.

Our findings speak to the ability of persuasive advertisement to help build trust that allows households to participate in equity markets. We also highlight the importance of local access to finance. Moreover, we document the complementary nature of local access and trust building through persuasive advertising. Together these findings suggest that under-served neighborhoods that lack access to investment advisors may not be able to reap the benefits of promotional campaigns that help encourage broader participation in the markets.

References

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey Wooldridge, 2017, When should you adjust standard errors for clustering?, National Bureau of Economic Research Working Paper.

Alesina, Alberto, and Eliana La Ferrara, 2002, Who trusts others?, *Journal of Public Economics* 85, 207–234.

Allen, Jason, Robert Clark, and Jean-François Houde, 2014, The effect of mergers in search markets: Evidence from the canadian mortgage industry, *American Economic Review* 104, 3365–96.

Allen, Jason, Robert Clark, and Jean-François Houde, 2019, Search frictions and market power in negotiated-price markets, *Journal of Political Economy* 127, 1550–1598.

Barber, Brad M, Terrance Odean, and Lu Zheng, 2005, Out of sight, out of mind: The effects of expenses on mutual fund flows, *Journal of Business* 78, 2095–2120.

Brown, James, J. Anthony Cookson, and Rawley Z. Heimer, 2019, Growing up without finance, *Journal of Financial Economics* 134, 591–616.

Brown, Jeffrey R, and Austan Goolsbee, 2002, Does the internet make markets more competitive? evidence from the life insurance industry, *Journal of political economy* 110, 481–507.

Brown, Jeffrey R, Zoran Ivković, Paul A Smith, and Scott Weisbenner, 2008, Neighbors matter: Causal community effects and stock market participation, *Journal of Finance* 63, 1509–1531.

Burke, Jeremy, and Angela A. Hung, 2015, Trust and financial advice, RAND Working Paper.

Christoffersen, Susan EK, and David K Musto, 2002, Demand curves and the pricing of money management, *The Review of Financial Studies* 15, 1499–1524.

Duarte, Fabian, and Justine S Hastings, 2012, Fettered consumers and sophisticated firms: evidence from mexico's privatized social security market, Technical report, National Bureau of Economic Research.

Elton, Edwin J, Martin J Gruber, and Jeffrey A Busse, 2011, Are investors rational? choices among index funds, in *Investments And Portfolio Performance*, 145–172 (World Scientific).

Gallaher, Steven T, Ron Kaniel, and Laura T Starks, 2015, Advertising and mutual funds: From families to individual funds, CEPR Discussion Paper No. DP10329.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny, 2015, Money doctors, *Journal of Finance* 70, 91–114.

Green, Richard C, Burton Hollifield, and Norman Schürhoff, 2007, Financial intermediation and the costs of trading in an opaque market, *The Review of Financial Studies* 20, 275–314.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2008, Trusting the stock market, *the Journal of Finance* 63, 2557–2600.

Gurun, Umit G, Gregor Matvos, and Amit Seru, 2016, Advertising expensive mortgages, *Journal of Finance* 71, 2371–2416.

Haliassos, Michael, and Carol C Bertaut, 1995, Why do so few hold stocks?, *the economic Journal* 105, 1110–1129.

Hilgert, Marianne A, Jeanne M Hogarth, and Sondra G Beverly, 2003, Household financial management: The connection between knowledge and behavior, *Fed. Res. Bull.* 89, 309.

Hong, Harrison, Jeffrey D Kubik, and Jeremy C Stein, 2004, Social interaction and stock-market participation, *Journal of Finance* 59, 137–163.

Hortaçsu, Ali, and Chad Syverson, 2004, Product differentiation, search costs, and competition in the mutual fund industry: A case study of s&p 500 index funds, *The Quarterly journal of economics* 119, 403–456.

Hung, Angela A., Noreen Clancy, Jeff Dominitz, Eric Talley, Claude Berribi, and Farrukh Suvankulov, 2008, Investor and industry perspectives on investment advisors and broker-dealers, Technical Report RAND Center for Corporate Ethics and Governance.

Jain, Prem C, and Joanna Shuang Wu, 2000, Truth in mutual fund advertising: Evidence on future performance and fund flows, *Journal of Finance* 55, 937–958.

Li, Xing, Wesley R Hartmann, and Tomomichi Amano, 2020, Preference externality estimators: A comparison of border approaches and IVs, Working Paper, Stanford University.

Lin, Leming, 2020, Bank deposits and the stock market, *The Review of Financial Studies* 33, 2622–2658.

Lusardi, Annamaria, and Olivia S. Mitchell, 2007, Financial literacy and retirement planning: New evidence from the RAND American Life Panel, MRRC Working Paper.

Mullainathan, Sendhil, Joshua Schwartstein, and Andrei Shleifer, 2008, Coarse thinking and persuasion, *Quarterly Journal of Economics* 123, 577–619.

Shapiro, Bradley T, 2018, Positive spillovers and free riding in advertising of prescription pharmaceuticals: The case of antidepressants, *Journal of Political Economy* 126, 381–437.

Shapiro, Bradley T, Günter J Hitsch, and Anna E Tuchman, 2021, TV advertising effectiveness and profitability: Generalizable results from 288 brands, Working paper, University of Chicago.

Spenkuch, Jörg L, and David Toniatti, 2018, Political advertising and election results, *The Quarterly Journal of Economics* 133, 1981–2036.

Tuchman, Anna E, 2019, Advertising and demand for addictive goods: The effects of e-cigarette advertising, *Marketing Science* 38, 994–1022.

van Rooij, Maarten, Annamaria Lusardi, and Rob Alessie, 2011, Financial literacy and stock market participation, *Journal of Financial Economics* 102, 449–472.

Vissing-Jorgensen, Annette, 2003, Perspectives on behavioral finance: Does "irrationality" disappear with wealth? evidence from expectations and actions, *NBER Macroeconomics Annual* 18, 139–194.

Figure 1: **Nielsen Designated Market Areas**

Figure 2: **Tallahassee-Jacksonville DMA Border**

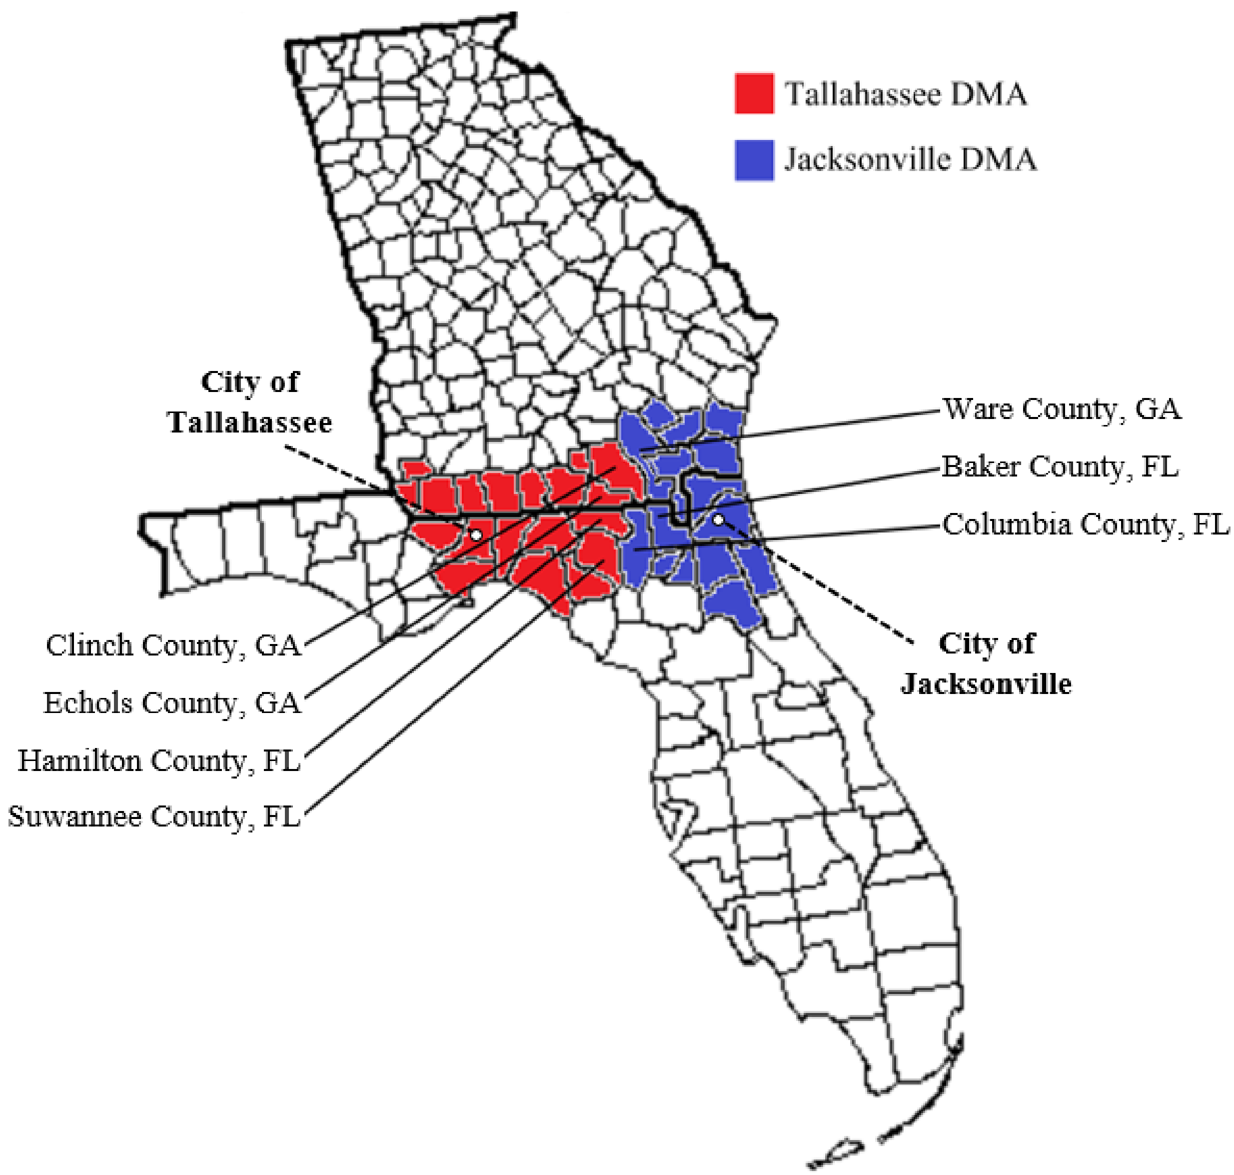

Figure 3: **Price of Spot TV Advertising vs. Financial Ad Units Purchased**

Panel A: Boston and New York City Markets

Panel B: Market Adjusted

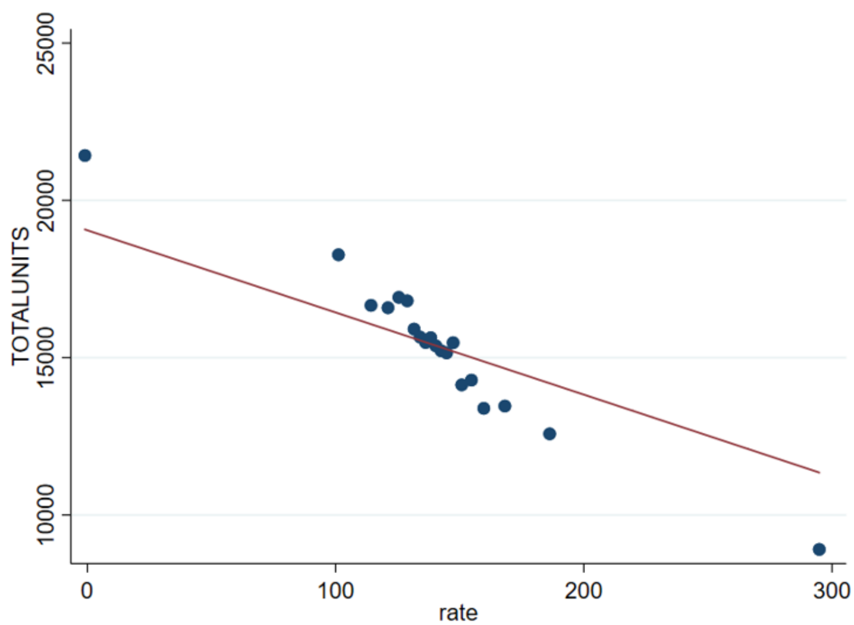

Table 1: **RIA Sample**

This table summarizes advertising and non-advertising RIAs from 2004 to 2018. The unit of observations is an RIA-year. RIA characteristics come from the advisor's Form ADV filings. An *Advertiser* is an RIA that matches to the Kantar Media database while a *Non-advertiser* is an RIA with no matching record to the Kantar Media advertising database. *AUM* is the RIA's assets under management in millions of US dollars. *Number of Accounts* is the RIA's total number of discretionary and non-discretionary accounts. *Employees* is the RIA's total number of full and part-time non-clerical employees. *Retail* is an indicator equal to 1 if 50% or more of the RIA's clients are individuals or high net worth individuals. *Planning* is an indicator equal to 1 if the RIA offers financial planning services. Standard errors (reported in parentheses) are clustered at the individual RIA level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Advertiser	Non-advertiser	Difference	t-stat
AUM (million of USD)	38,363	3,010	35,353	5.38***
Number of Accounts	1,748	392	1,356	8.38***
Employees	161	31	130	12.55***
Retail	54.8%	48.5%	6.25%	3.63***
Offer Planning Services	46.1%	36.7%	9.5%	5.19***

Table 2: **Kantar Financial Advertising Sample**

This table summarizes the advertising expenditure and financial advisor characteristics for the small border counties sample from 2004 to 2018. Small borders are defined as borders where the combined population of all counties on each side of the DMA border is less than 10% of the total population in its respective DMA. *Total Ad Units* is the RIAs' local advertising units in the current year. *Total Ad Dollars* is the RIAs' local advertising expenditure in the current year. *Ad $ per household* is defined as the RIAs' national advertising expenditure scaled by the national population plus their local advertising expenditure scaled by the population of the DMA. *# of RIAs* is the number of RIAs with a physical presence in the county-year.

	Mean	Std. Dev	25^{th}	50^{th}	75^{th}
Total Ad Units	31,008.3	49,171.8	2,371	11,358	40,771
Total Ad Dollars	$ 7,734,262	18,900,000	148,797	1,445,481	7,414,517
Ad $ per household	$ 26.53	5.96	23.08	25.58	28.92
# of RIAs	6.4	11.3	1	3	8

Table 3: **IRS Summary Statistics**

This table reports the summary statistics on the number of tax return filers in 2018 and stock market participation rates by adjusted gross income. The stock market participation rate is defined as the fraction of household tax returns that report dividend income. Panel A reports the summary statistics in the IRS's Statistics of Income sample. Panel B reports the summary statistics in our small border counties sample from 2004 to 2018. Small borders are defined as borders where the combined population of all counties on each side of the DMA border is less than 10% of the total population in its respective DMA.

Panel A: Full United States Sample

Income	Households (millions, 2018)	Participation
$1-25k	48.9	10.9%
25 - 50k	35.7	15.7%
50 - 75k	20.9	23.1%
75 - 100k	13.6	27.8%
100 - 200k	20.8	35.9%
200k+	8.2	56.6%
U.S.	148.2	24.8%

Panel B: Small Borders Sample

Income	Households (millions, 2018)	Participation
$1-25k	5.0	9.3%
25 - 50k	3.5	14.2%
50 - 75k	2.0	21.8%
75 - 100k	1.3	26.8%
100 - 200k	1.8	36.6%
200k+	0.5	52.9%
U.S.	14.2	25.0%

Table 4: **Baseline**

This table reports estimates from regressions of stock market participation rates on RIAs' local advertising spending. We restrict our sample to the small border counties from 2004 to 2018. Small borders are defined as borders where the combined population of all counties on each side of the DMA border is less than 10% of the total population in its respective DMA. The unit of observation is the county-year. We create a county-year level proxy of stock market participation rate based on the measure from Brown et al. (2008). The county-year level stock market participation rate is defined as the fraction of household tax returns that report dividend income. *Log (Units of Advertising)* is the logarithm of RIAs' local advertising units in the current year. *Log ($ of Advertising)* is the logarithm of RIAs' local advertising expenditure in the current year. *$ per household* is defined as the RIAs' national advertising expenditure scaled by the national population plus their local advertising expenditure scaled by the population of the DMA. All specifications include Border × DMA and Border × Year fixed effects. Standard errors (reported in parentheses) are clustered at the DMA level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)
Log(Units of Advertising)	0.0016*		
	(0.0008)		
Log($ of Advertising)		0.0011**	
		(0.0005)	
$ per household			0.0017**
			(0.0007)
$ per household squared			-0.00003***
			(0.00001)
Border × DMA FE	Yes	Yes	Yes
Border × Year FE	Yes	Yes	Yes
R^2	0.679	0.679	0.679
Observations	19,086	19,099	19,099

Table 5: **Effects by Household Income Level**

This table reports estimates from regressions of stock market participation rates on RIAs' local advertising spending, interacted with different AGI tier. We restrict our sample to small border counties from 2004 to 2018. Small borders are defined as borders where the combined population of all counties on each side of the DMA border is less than 10% of the total population in its respective DMA. The unit of observation is the county-AGI tier-year. The county-year level stock market participation rate is defined as the fraction of household tax returns that report dividend income. *Log(Units of Advertising)* is the logarithm of RIAs' local advertising units in the current year. *Log($ of Advertising)* is the logarithm of RIAs' local advertising expenditure in the current year. We include the following household AGI tiers: 1-75k, 75k-100k, 100k-200k, and >200k. All specifications include Border × DMA, Border × Year, and AGI Tier × Year fixed effects. Standard errors (reported in parentheses) are clustered at the DMA and year. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)
200k+ × Log(Units of Advertising)	0.0123***	
	(0.0038)	
100k-200k × Log(Units of Advertising)	0.0042**	
	(0.0018)	
75k-100k × Log(Units of Advertising)	-0.0004	
	(0.0015)	
Under 75k × Log(Units of Advertising)	-0.0017	
	(0.0014)	
200k+ × Log($ of Advertising)		0.0105***
		(0.0029)
100k-200k × Log($ of Advertising)		0.0026*
		(0.0014)
75k-100k × Log($ of Advertising)		-0.0009
		(0.0010)
Under 75k × Log($ of Advertising)		-0.0017
		(0.0013)
Border × DMA FE	Yes	Yes
Border × Year FE	Yes	Yes
AGI Tier × Year FE	Yes	Yes
R^2	0.695	0.695
Observations	107,830	107,939

Table 6: **Placebo Tests**

This table reports placebo test results. We restrict our sample to the small border counties from 2004 to 2018. Small borders are defined as borders where the combined population of all counties on each side of the DMA border is less than 10% of the total population in its respective DMA. The unit of observation is the county-year. In columns (1) and (2), the dependent variable is the average household income level. *Log(Units of RIA Advertising)* is the logarithm of RIAs' local advertising units in the current year. *Log($ of RIA Advertising)* is the logarithm of RIAs' local advertising expenditure in the current year. In column (3) – (5), the dependent variable is the county-year level stock market participation rate. *Log(Units of Credit Union Advertising)* is the logarithm of credit union firms' local advertising units in the current year. *Log(Units of Loan Advertising)* is the logarithm of loan institutions' local advertising units in the current year. *Log(Units of Lending Advertising)* is the logarithm of lending firms' local advertising units in the current year. All specifications include Border × DMA and Border × Year fixed effects. Standard errors (reported in parentheses) are clustered at the DMA level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Avg. Salary		Participation		
	(1)	(2)	(3)	(4)	(5)
Log(Units of RIA Advertising)	-0.0522				
	(0.0450)				
Log($ of RIA Advertising)		-0.0392			
		(0.0573)			
Log(Units of Credit Union Advertising)			0.0002		
			(0.0004)		
Log(Units of Loan Advertising)				0.0004	
				(0.0005)	
Log(Units of Lending Advertising)					-0.0003
					(0.0003)
Border × DMA FE	Yes	Yes	Yes	Yes	Yes
Border × Year FE	Yes	Yes	Yes	Yes	Yes
R^2	0.754	0.754	0.688	0.673	0.683
Observations	19,086	19,099	17,933	17,267	10,219

Table 7: **Local Presence and Advertising**

This table reports estimates from regressions of stock market participation rates on the presence of local RIAs and local advertising spending by RIAs. Panel A reports the average effects across all households and Panel B reports the effects by income levels. We restrict our sample to the small border counties from 2004 to 2018. Small borders are defined as borders where the combined population of all counties on each side of the DMA border is less than 10% of the total population in its respective DMA. The county-year level stock market participation rate is defined as the fraction of household tax returns that report dividend income. *Any Local RIA* is an indicator variable equal to one if any RIA has physical presence in the county-year. *Any Local Advertising RIA* is an indicator variable equal to one if any advertising RIA has physical presence in the county-year. *Log(Units of Local RIA Advertising)* is the logarithm of the units of local advertising purchased by an RIA with a local presence in the county-year. *Log(Units of Non-local RIA Advertising)* is the logarithm of the units of local advertising purchased by an RIA without a local presence in the county-year. *Log ($ of Local RIA Advertising)* is the logarithm of the dollar amount of local advertising purchased by an RIA with a local presence in the county-year. *Log ($ of Non-local RIA Advertising)* is the logarithm of the dollar amount of local advertising purchased by an RIA without a local presence in the county-year. *Test of difference: Local Units vs. Non-Local Units* reports the test of the difference in the coefficients on *Log(Units of Local RIA Advertising)* and *Log(Units of Non-local RIA Advertising)* in column (2). *Test of difference: Local $ vs. Non-Local $* reports the test of the difference in the coefficients on *Log ($ of Local RIA Advertising)* and *Log ($ of Non-local RIA Advertising)* in column (3). All specifications include border × DMA and border × year fixed effects. Standard errors (reported in parentheses) are clustered at the DMA level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Average Effects

	(1)	(2)	(3)
Any Local RIA	-0.0012	-0.0000	-0.0016
	(0.0052)	(0.0054)	(0.0054)
Any Local Advertising RIA	0.0176***		
	(0.0044)		
Log(Units of Local RIA Advertising)		0.0032***	
		(0.0008)	
Log(Units of Non-local RIA Advertising)		0.0007	
		(0.0010)	
Log($ of Local RIA Advertising)			0.0020***
			(0.0005)
Log($ of Non-local RIA Advertising)			0.0003
			(0.0008)
Border×DMA FE	Yes	Yes	Yes
Border×Year FE	Yes	Yes	Yes
R^2	0.689	0.687	0.689
Observations	19,764	19,086	19,099
Test of difference: Local Units vs. Non-Local Units		0.0025*	
Test of difference: Local $ vs. Non-Local $			0.0017*

Panel B: Effects by Income Levels

	(1)	(2)
Any Local RIA	0.0169*	0.0142*
	(0.0081)	(0.0078)
200k+× Log(Units of Local RIA Advertising)	0.0237***	
	(0.0037)	
100k-200k× Log(Units of Local RIA Advertising)	0.0059***	
	(0.0013)	
75k-100k× Log(Units of Local RIA Advertising)	0.0058***	
	(0.0015)	
Under 75k× Log(Units of Local RIA Advertising)	0.0007	
	(0.0008)	
200k+× Log(Units of Non-local RIA Advertising)	-0.0040	
	(0.0031)	
100k-200k× Log(Units of Non-local RIA Advertising)	0.0022	
	(0.0018)	
75k-100k× Log(Units of Non-local RIA Advertising)	-0.0022	
	(0.0014)	
Under 75k× Log(Units of Non-local RIA Advertising)	-0.0001	
	(0.0010)	
200k+× Log($ of Local RIA Advertising)		0.0162***
		(0.0026)
100k-200k× Log($ of Local RIA Advertising)		0.0043***
		(0.0008)
75k-100k× Log($ of Local RIA Advertising)		0.0036***
		(0.0009)
Under 75k× Log($ of Local RIA Advertising)		0.0005
		(0.0005)
200k+× Log($ of Non-local RIA Advertising)		-0.0028
		(0.0022)
100k-200k× Log($ of Non-local RIA Advertising)		0.0010
		(0.0012)
75k-100k× Log($ of Non-local RIA Advertising)		-0.0016
		(0.0009)
Under 75k× Log($ of Non-local RIA Advertising)		0.0005
		(0.0008)
Border×DMA FE	Yes	Yes
Border×Year FE	Yes	Yes
AGI Tier×Year FE	Yes	Yes
R^2	0.710	0.713
Observations	107,830	107,939

Table 8: **County-level Heterogeneous Effects: Stock Market Participation Rates for 100k+ Households**

This table reports estimates from regressions of stock market participation rates for households with annual income in excess of $100,000 on the presence of local RIAs and local advertising RIAs, interacted with county-level income and racial disparity measures. We restrict our sample to the small border counties from 2004 to 2018. Small borders are defined as borders where the combined population of all counties on each side of the DMA border is less than 10% of the total population in its respective DMA. The county-year level stock market participation rate is defined as the fraction of household tax returns that report dividend income. *Any Local RIA* is an indicator variable equal to one if any RIA has physical presence in the county-year. *Any Local Advertising RIA* is an indicator variable equal to one if any advertising RIA has physical presence in the county-year. *High Income Disparity* is an indicator variable equal to one if the county-year level normalized Herfindahl index based on household total income is above the median for all counties in the given year. *High Racial Disparity* is an indicator variable equal to one if the county-year level normalized Herfindahl index based on ratios of all race groups is above the median for all counties in the given year. All specifications include border × DMA and border × year fixed effects. Standard errors (reported in parentheses) are clustered at the DMA level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)
Any Local RIA	0.0112	0.0408**
	(0.0136)	(0.0184)
Any Local Advertising RIA	0.0206**	0.0379***
	(0.0081)	(0.0117)
High Income Disparity× Any Local Advertising RIA	0.0353***	
	(0.0113)	
High Racial Disparity× Any Local Advertising RIA		0.0257*
		(0.0140)
High Income Disparity× Any Local RIA	0.0303*	
	(0.0172)	
High Racial Disparity× Any Local RIA		0.0139
		(0.0260)
High Income Disparity	-0.0927***	
	(0.0169)	
High Racial Disparity		-0.0427*
		(0.0241)
Racial Group Control	–	Yes
Border×DMA FE	Yes	Yes
Border×Year FE	Yes	Yes
R^2	0.635	0.641
Observations	16,518	16,528

Appendices

Table A1: **Test for Selection Across Borders**

This table reports the average number of financial advisors (scaled by the population), income, population, education attainment, debt-to-income ratios, social capital index, and age on both the sides of a DMA border. We calculate average values across counties and aggregate them to the border-DMA level. We consider the side of the border with greater (less) advertising to be the Heavy (Light) side of a DMA border. ***,**, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Heavy	**Light**	**t-stat**	**p-value**
Advisors per 10,000	10.06	10.55	-0.50	0.616
	(0.60)	(0.76)		
Income	31,609.77	30,826.50	1.22	0.223
	(436.47)	(470.04)		
Population	52,216	66,714	-1.20	0.231
	(4,038)	(11,390)		
% College Degree	13.54	13.61	-0.15	0.883
	(0.33)	(0.35)		
Debt to Income	1.64	1.66	-0.30	0.763
	(0.06)	(0.07)		
Social Capital Index	0.027	-0.008	0.03	0.746
	(0.08)	(0.08)		
Age	40.02	40.12	-0.35	0.729
	(0.20)	(0.20)		

Table A2: **Retail RIA Advertising**

This table reports estimates from regressions of stock market participation rates on local advertising spending by retail RIAs. We restrict our sample to the small border counties from 2004 to 2018. Small borders are defined as borders where the combined population of all counties on each side of the DMA border is less than 10% of the total population in its respective DMA. The county-year level stock market participation rate is defined as the fraction of household tax returns that report dividend income. *Log(Units of Retail RIA Advertising)* is the logarithm of the units of local advertising purchased by an RIA with more than 50% individual clients or more than 50% high net worth clients in the county-year. *Log(Units of Non-retail RIA Advertising)* is the logarithm of the units of local advertising purchased by an RIA with less than or equal to 50% individual clients and less than or equal to 50% high net worth clients in the county-year. *Log ($ of Retail RIA Advertising)* is the logarithm of the dollar amount of local advertising purchased by an RIA with more than 50% individual clients or more than 50% high net worth clients in the county-year. *Log ($ of Non-retail RIA Advertising)* is the logarithm of the dollar amount of local advertising purchased by an RIA with less than or equal to 50% individual clients and less than or equal to 50% high net worth clients in the county-year. *Test of difference: Local Units vs. Non-Local Units* reports the test of the difference in the coefficients on *Log(Units of Retail RIA Advertising)* and *Log(Units of Non-retail RIA Advertising)* in column (1). *Test of difference: Retail $ vs. Non-retail $* reports the the test of the difference in the coefficients on *Log ($ of Retail RIA Advertising)* and *Log ($ of Non-retail RIA Advertising)* in column (2). All specifications include border × DMA and border × year fixed effects. Standard errors (reported in parentheses) are clustered at the DMA level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)
Log(Units of Retail RIA Advertising)	0.0020**	
	(0.0009)	
Log(Units of Non-retail RIA Advertising)	0.0002	
	(0.0005)	
Log($ of Retail RIA Advertising)		0.0016**
		(0.0007)
Log($ of Non-retail RIA Advertising)		-0.0003
		(0.0003)
Border×DMA FE	Yes	Yes
Border×Year FE	Yes	Yes
R^2	0.676	0.676
Observations	18,665	18,652
Test of difference: Retail Units vs. Non-retail Units	0.0018**	
Test of difference: Retail $ vs. Non-retail $		0.0019**

Table A3: **Alternative Specifications**

This table reports estimates from regressions of stock market participation rates on financial advisers' advertising spending with alternative specifications. In column (1), we report estimates in the small border counties sample using TV advertising spending. The independent variable in column (1) is the logarithm of financial advisers' local TV advertising units in the current year. We include border $\times$ DMA and border $\times$ year fixed effects in column (1). In column (2), we report estimates in all counties sample. The independent variable in column (2) is the logarithm of financial advisers' local advertising units in the current year. We include county and year fixed effects in column (2). Column (3) reports estimates in a sample of county pairs that are each less than 3% of their DMA population. The independent variable in column (3) is the logarithm of financial advisers' local advertising units in the current year. We include county $\times$ pair and pair $\times$ year fixed effects in column (3). Standard errors (reported in parentheses) in all specifications are clustered at the DMA level. ***,**, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	TV Only (1)	Panel (2)	County Pair Approach (3)
Log(Units of Advertising)	0.0009**	0.0016***	0.0017**
	(0.0004)	(0.0004)	(0.0008)
Border×DMA FE	Yes		
Border×Year FE	Yes		
County FE		Yes	
Year FE		Yes	
County×Pair FE			Yes
Pair×Year FE			Yes
R^2	0.682	0.883	0.925
Observations	19,764	41,194	24,508