The Promises and Pitfalls of Robo-advising

Francesco D'Acunto, Nagpurnanand Prabhala, Alberto G. Rossi

CESifo

The Promises and Pitfalls of Robo-advising

Abstract

We study a robo-advising portfolio optimizer that constructs tailored strategies based on investors' holdings and preferences. Adopters are similar to non-adopters in terms of demographics, but have more assets under management, trade more, and have higher risk-adjusted performance. The robo-advising tool has opposite effects across investors with different levels of diversification before adoption. It increases portfolio diversification and decreases volatility for those that held less than 5 stocks before adoption. These investors' portfolios perform better after using the tool. At the same time, robo-advising barely affects diversification for investors that held more than 10 stocks before adoption. These investors trade more after adoption with no effect on average performance. For all investors, robo-advising reduces - but does not fully eliminate - pervasive behavioral biases such as the disposition effect, trend chasing, and the rank effect, and increases attention based on online account logins. Our results emphasize the promises and pitfalls of robo-advising tools, which are becoming ubiquitous all over the world.

Francesco D'Acunto
R.H. Smith School of Business
University of Maryland
College Park / MD / USA
fdacunto@rhsmith.umd.edu

Nagpurnanand Prabhala
R.H. Smith School of Business
University of Maryland
College Park / MD / USA
nprabhala@rhsmith.umd.edu

Alberto G. Rossi
R.H. Smith School of Business
University of Maryland
College Park / MD / USA
arossi@rhsmith.umd.edu

1 Introduction

Most investors would benefit from stock market participation (Campbell and Viceira, 2002; Campbell, 2006). The benefits of participation depend on whether investors hold diversified portfolios. In practice, investors do not diversify (Badarinza, Campbell, and Ramadorai, 2016). Financial advising can potentially mitigate under-diversification by helping investors move towards more diversified portfolios, but financial advisers are prone to behavioral biases and display cognitive limitations (Linnainmaa, Melzer, and Previtero, 2017). Our study focuses on a FinTech robo-advising tool that delivers diversification advice to individual investors and does not require the intervention of human advisers. We examine the uptake of the tool and assess its impact on financial decision-making by investors. We find that robo-advising has opposite effects on investors' performance based on their level of sophistication, whereas it reduces a set of well-known behavioral biases for all investors.

The robo-advising tool we examine is an automated portfolio optimizer introduced by a brokerage firm to its clients in India. The tool uses Markowitz mean-variance optimization to construct optimal portfolio weights based on historical data and modern techniques such as shrinkage and short-selling constraints. The tool is flexible as it allows investors to rebalance current portfolios and add extra stocks from a set of up to 15 liquid stocks the brokerage house chooses. Importantly, the tool incorporates simplified trade execution. Investors merely need to click a button to execute all necessary trades in batch mode.

We interpret the robo-adviser as a way to simplify the set of decisions investors have to make to rebalance their portfolios. When investors have no access to the tool, rebalancing involves a complex set of decisions. Investors face the daunting task of choosing from a large number of securities and allocating their wealth among the chosen stocks. To simplify this set of problems, investors often use suboptimal rules of thumb (e.g., Frydman, Hartzmark, and Solomon, Forthcoming). Robo-advising simplifies the process, because automated execution lets investors implement easily the advice they receive.

We investigate three dimensions of robo-advising, that is, the take up of the optimizer, its effect on portfolio outcomes such as volatility and performance, and on the incidence of behavioral biases. We report single-difference results that control for time-invariant investor characteristics, as well as difference-in-differences results that exploit quasi-random variation in the introduction of the robo-

2

advising tool. Our dataset contains information on investors' demographic characteristics as well as their trading histories, portfolio holdings, and access to both human advice and robo-advice.

We first analyze the determinants of adopting the robo-advising portfolio optimizer, which sheds light on the types of investors that are more receptive to technological innovation in the realm of financial advice. We find that users and non-users are indistinguishable along several demographic characteristics, including their gender, age, and trading experience. At the same time, users have a larger amount of wealth invested with the brokerage house and are more sophisticated. They are more involved with the management of their portfolios and have superior risk-adjusted performance, consistent with Gargano and Rossi (2017).

We next analyze the effects of robo-advising on portfolio diversification, risk, and investment returns in a within-investor analysis, which partials out all the time-invariant determinants of adoption. Investors holding less than 10 stocks before using the optimizer increase the number of stocks they hold and experience sharp declines in portfolio volatility. For investors with 10 or more stocks, the number of stocks held decreases after portfolio optimizer usage, suggesting that the optimizer recommends closing positions in stocks that would be shorted had the short-sales constraint not been binding. While these investors hold fewer stocks after adoption, portfolio volatility does not increase but decreases less compared to undiversified individuals. The evidence that undiversified investors benefit more from robo-advice – whose technology makes implementation of advice simple also for the less savvy investors – suggests that robo-advice can be an effective tool to help investors diversify their portfolios, compared to other forms of advice (see, Bhattacharya et al., 2012, Linnainmaa, Melzer, and Previtero, 2017).

We move on to assess the effects of the usage of the portfolio optimizer on post-adoption trading. Once again, we sort investors based on their levels of diversification before usage. We find that market-adjusted investment performance improves for less diversified investors. The average returns for the ex-ante diversified investors are essentially flat. These investors pay more attention to their portfolio and increase their trading volume, which we proxy by the overall amount of trading fees. In line with the results described above, these findings suggest that the robo-advising tool conveys more benefits to investors who are less diversified ex-ante.

Our third set of tests examine prominent behavioral biases individual investors exhibit when buying and selling stocks. On the one hand, the trades suggested by the robo-advising tool should not reflect

any behavioral biases.[1] A correction of behavioral biases could therefore be mechanical or could stem from the fact that investors learn how to place unbiased trades as they follow the robo-advising tool. On the other hand, because investors trade more after using the robo-advising tool, robo-advice could exacerbate the effects of behavioral biases if investors increased the number of trades they placed independently.

For selling decisions, we examine the disposition effect, whereby investors are more likely to realize gains than losses on their positions (Shefrin and Statman, 1985). To assess buying behavior, we examine trend chasing, whereby investors tend to purchase stocks after a set of positive returns (Barber and Odean, 2008). We also examine the rank effect, whereby investors are more likely to trade the best and worst performing stocks in their portfolios (Hartzmark, 2014). We test the incidence of all three biases before and after investors access robo-advice. The biases are substantially less pronounced after the usage of the portfolio optimizer regardless of investors' diversification before usage, even if the tool does not fully eliminate them.

The results described above are based on single-difference tests, in which we compare diversification, trading behavior, and trading performance within individuals, before and after usage of the portfolio optimizer. The single-difference tests exclude that our results are driven by systematic, time-invariant investor characteristics, and hence by the selection into usage of the portfolio optimizer. At the same time, the single-difference tests do not allow us to rule out that time-varying shocks to trading motives cause both the usage of the optimizer and the change in trading behavior after usage.

To address these identification concerns, we propose a strategy that exploits the quasi-random variation in adoption induced by the way the portfolio optimizer is introduced to the market. We build on the fact that at several points in time, the brokerage house asks human advisers to call their clients to promote the usage of the portfolio optimizer and initiate usage of the tool. The brokerage house had no underlying motivations for pushing the portfolio optimizer at any point in time, apart from the fact that their technology team thought the device was ready to use broadly and they wanted to market it as a free service to their clients. Crucially for our purposes, our dataset identifies all the outbound and inbound calls human advisers have with clients at each point in time. Moreover, we

---

[1]Recent research suggests human advisors might themselves be subject to behavioral biases, and hence transmit such biases to the trading behavior of their clients (see Linnainmaa, Melzer, and Previtero, 2017) Because robo-advising algorithms are designed by humans, these algorithms might themselves reflect the behavioral biases of those designing them.

know whether calls went through and, if yes, the call length.

In our identification strategy, treated clients are those clients the human advisers reach in the days in which they are promoting the portfolio optimizer, and use the optimizer during the call. Control clients are those clients human advisers *try* to contact on the same day to promote the optimizer, but do not answer the phone, and hence are not exposed to the tool.[2] The subset of clients advisers call might not be chosen at random. Advisers might call clients whose characteristics make them more likely to adopt the optimizer, or clients they think would benefit the most from using the optimizer. But this potential selection is not a problem – and perhaps even advantageous – for our difference-in-differences strategy. In the control sample are clients who do not answer the phone, despite being as likely to benefit from the optimizer as clients that answer the phone. Overall, the difference-in-differences specifications confirm our results.

2   Related Literature

Our work contributes to multiple strands of literature in Finance and Economics. First, we contribute to the literature in household finance. Campbell (2006) points out in his presidential address that the benefits of financial markets depend on how effectively households use financial products.[3]

Participation in the stock market is optimal from a portfolio allocation viewpoint given the historically high risk premia of stock market investments. However, attaining these high returns depends on whether investors hold appropriately diversified portfolios. The actual risky holdings of investors deviate considerably from theoretical predictions (Badarinza, Campbell, and Ramadorai, 2016). In particular, participants in the stock market tend to be under-diversified. Undiversified portfolios result in investors bearing idiosyncratic risk that is not compensated by higher returns.

Financial advising can potentially help mitigate underdiversification and help investors realize better outcomes (Gennaioli, Shleifer, and Vishny, 2015). However, for many retail investors, traditional financial advisers are too costly. FinTech robo-advising makes access to the financial advice available

---

[2]We require that non-responsive clients do not use the portfolio optimizer in the thirty days after the attempted call by their human adviser. The results are not sensitive to using different horizons for this restriction.

[3]Recent work in this area addresses practical questions on the design or delivery of financial services and also informs policies such as those on tax, investor protection, financial literacy, or investor education. See, e.g., Anagol, Balasubramaniam, and Ramadorai (2017), Barber and Odean (2000, 2008), Barberis and Thaler (2003), Calvet, Campbell, and Sodini (2009), Grinblatt and Keloharju (2001a,b) for evidence on investor behavior.

at low cost.[4] Moreover, advisers often adopt a one-size-fits-all approach and might be prone to behavioral biases or display cognitive limitations (Linnainmaa, Melzer, and Previtero, 2017). Robo-advising tools might be subject to the biases, conflicts, and limitations of the humans and institutions that develop them. However, robo-advising is by construction neutral to the idiosyncrasies of specific human advisers.

Our study is also relevant to the broader literature in Economics on technology adoption (see Comin and Mestieri, 2014, for a recent review). New technology and its adoption play an important role in improving productivity and economic growth (Romer, 1990; Aghion and Howitt, 1992). Comin and Mestieri (2014) point out that a key gap in this literature is the lack of micro-level datasets on adoption. This is an important issue because technological progress is in large part due to the adoption of new technologies, not just their creation. Our study contributes to this literature by analyzing granular, micro-data on the adoption of a particular technology. Our data allow us to estimate both the intended and unintended consequences of adopting this technology on measurable outcomes at the portfolio level as well as on the incidence of behavioral biases among investors.

Within technology adoption, we are among the first papers to study the effects of FinTech on individual outcomes. With few exceptions (e.g., Tufano, 1989), this is an area that has seen relatively little research.[5] The scarcity of work on technological innovations in finance lead Frame and White (2004) to write that "... Everybody talks about financial innovation, but (almost) nobody empirically tests hypotheses about it" in reference to a quote attributed to Mark Twain. Since the remark by Frame and White (2004), there has been substantial work in Development Economics on studying the introduction and evaluation of new financial products aimed at the poor, which are typically unbanked individuals unfamiliar with relatively well-known financial products (e.g., see Dupas and Robinson, 2013). Little work exists on the effects of financial technology aimed at the investment decisions of high-income households. We contribute towards filling this gap.

---

[4]See, e.g., the Ernst and Young report "Advice goes virtual." An S&P global report predicts that by 2021, robo-advising will have assets under management of $450 billion (https://marketintelligence.spglobal.com/our-thinking/ideas/u-s-digital-adviser-forecast-aum-to-surpass-450b-by-2021, accessed October 11, 2017.)

[5]Other work focuses on agriculture (Conley and Udry, 2010; Bold et al., Forthcoming), health products (Dupas, 2014), or manufacturing (Atkin et al., 2015). Manuelli and Seshadri (2014) analyze technological adoptions in the tractor industry between 1910 and 1960, while Skinner and Staiger (2015) and Chandra et al. (2016) study the role of innovation on the health care industry using Medicare data.

3 Robo-advising

Robo-advising is the delivery and execution of financial advice through automated algorithms on digital platforms. We study robo-advisers that assist individual investors in portfolio selection. Robo-advisers that perform these tasks have many similarities although there are some variations in the details across different platforms.[6] We briefly discuss the features of the major robo-advising services.

3.1 Robo-advising Industry: An Overview

Robo-advising is a growing industry. The 2016 S&P Global Market Intelligence Report puts the robo-advised assets under management at \$98.62 billion, and projects them to grow by over 40% annually.

The robo-advising industry started as a disruptive play by new entrants. These providers saw robo-advising as an opportunity to help customers human advisers would not serve. As of December 2017, the largest among the new players include Betterment, Wealthfront, and Personal Capital. Their 2017 Form ADVs indicate that the firms have assets under management (AUM) of about \$12 billion, \$9 billion, and \$4 billion, respectively, from about 306,000, 171,000, and 11,300 customers. Far less than 10% of these customers are high-net-worth individuals. Each of the firms employs between 100 and 200 employees and offers portfolio management services to clients. Recognizing opportunities in this market, such as including banks, fund houses, and brokerage firms are responding with hybrid forms of human- and robo-advising for all types of clients.

Robo-advisers might represent a significant improvement over human financial advisers for a number of reasons. First, robo-advisers typically use replicable algorithms based on financial theory. Second, robo-advisers employ technology that in many instances can simplify and speed up contact with clients. For example, robo-advisers can push out alerts to clients quickly in response to news or market changes.

Robo-advisers are also transparent. The interaction between human advisers and clients, on the other hand, often resembles a sales transaction, in which the adviser has an incentive to maximize personal incentives that may differ from investors' first best interests. The conflicts, biases, and

---

[6]See, e.g., the March 2016 FINRA "Report on Digital Investment Advice," available at `https://www.finra.org/sites/default/files/digital-investment-advice-report.pdf`, accessed May 1, 2017.

cognitive limitations of advisers can be transmitted to clients (e.g., Mullainathan, Noeth, and Schoar, 2012; Linnainmaa, Melzer, and Previtero, 2017). This point opened a policy debate on the costs and benefits of the so-called fiduciary rule that requires financial advisers to act in the best interests of the client. The 2015 edition of the "Investor Pulse" survey conducted by BlackRock is consistent with these observations. Top reasons for picking robo-advisers include convenience, simplicity, and not being pushed into products customers think they do not need.

An additional advantage of robo-advisers is the greater simplicity and efficiency in implementing strategies due to built-in automated algorithms. For example, tax-loss harvesting, a key feature offered by robo-advisers such as Wealthfront and Betterment, is facilitated by automated execution.

Digital advising is not necessarily without pitfalls. The ability to experiment with portfolio choices and the requirement that the investor ultimately authorizes any trade confers the investor complete control over the investment process. Giving investors such control is useful because it can help overcome algorithm aversion (Dietvorst, Simmons, and Massey, Forthcoming). At the same time, full control puts into play potential suboptimal behavior due to the lack of self control (Thaler and Shefrin, 1981) and might induce overtrading (Barber and Odean, 2000). Some robo-advisers have also faced criticism because they may put company profits ahead of investors' interests.[7]

Several economic incentives might drive both incumbents and new entrants to design robo-advising tools. As for incumbents, robo-advising allows reducing the costs of maintaining a full floor of human financial advisers. Human advisers are costly not only in terms of salary expenses, but also because they have high turnover that requires the firm to engage in significant training expenses on an ongoing basis.

Moreover, both incumbent and new entrants expect the overall market for financial advice to expand tremendously over the coming years. Being ahead of the competition in the robo-advising space is crucial to acquire the largest possible share of the new customers entering this market.

Incumbents and new entrants might follow different strategies when entering the robo-advisory space, based on economic incentives. Incumbent players have an existing base of customers to whom they need to cater. From a marketing perspective, introducing a disruptive product is difficult for them, because such product might debase their existing customers. Incumbents are also likely to have skilled employees they can deploy on the new products. As a result, they tend to offer additional

---

[7]See, for example, "Should Retirees Use Robo Advisers?, *Wall Street Journal*, November 12, 2017.

services to their clients, such as hybrid forms of human- and robo-advising. For instance, the company with which we work introduced a robo-advising tool in addition to the existing human advisory services that their customers had been using for year.

New entrants have more flexibility when it comes to introducing innovative products, because their strategy is acquiring market share by attracting new customers to financial advice or attracting existing customers incumbents serve. New entrants are thus more likely to propose financial products that have lower fees but offer fewer support services to their clients, which would require an infrastructure they do not have in place.

3.2   Implementing Robo-advising

The fundamental building block of robo-advisers is the classical Markowitz (1952) mean-variance optimization. This approach takes as inputs the vector of mean returns and the variance-covariance matrix and returns a set of efficient portfolios.

Despite its undeniable influence on the asset management industry, mean-variance optimization has a number of limitations. As a one-period model, mean-variance optimization does not consider time variation in the investment opportunity set. Neither does it consider explicitly that the efficient frontier is a function of each individual investor's horizon. The framework also assumes that returns are jointly normally distributed, while substantial empirical evidence shows returns are significantly fat-tailed.

Implementation also faces several challenges. A key difficulty is getting a precise estimate of the variance-covariance matrix. A standard way to reduce the effect of estimation error is to use shrinkage (e.g., see Ledoit and Wolf, 2004) or Bayesian techniques (e.g., Black and Litterman, 1991). The exact implementation details are typically kept proprietary. For example, Wealthfront simply notes that it uses both historical stock-market data and options data to infer volatility while Betterment indicates that it modifies the Black-Litterman approach. The precise assets from which the portfolios are drawn can vary. Schwab considers US and international equities, US and international treasuries, corporate bonds, TIPS, municipal bonds, and gold. Wealthfront and Betterment have narrower focus on US stocks and bonds. Investment strategies are usually implemented using ETFs, which are liquid and can be traded at low costs. Interestingly, Wealthfront offers a direct indexing product that invests in

individual stocks, as it leads to more efficient tax-loss harvesting.

3.3 The Robo-advising Tool We Study

The robo-advising technology we study – named *Portfolio Optimizer* – focuses on equities. Investors can access the portfolio optimizer from their online accounts. While investors have the option to enter the tickers they wish to consider in their portfolio allocation, the portfolio optimizer by default loads the investors' stock portfolio directly from their account. This feature of the optimizer aims at simplifying investors' access to the tool.

By default, the optimizer maximizes the investor's Sharpe ratio. The investor also has the option to specify the expected risk or return of the portfolio, but this happens in less than 5% of the cases. The application proposes the optimal portfolio weights according to Markowitz mean-variance optimization. To estimate the variance-covariance matrix, the algorithm uses 3 years of historical daily observations. To limit the effects of estimation error and to guarantee well-behaved portfolio weights, the algorithm implements modern techniques, such as shrinkage of the variance-covariance matrix. Moreover, the tool imposes short-sale constraints. Finally, investors need not contribute additional financial resources to their brokerage account to transition to the recommended portfolio. All these details of the computation of the optimal portfolio weights are accessible to investors. The robo-adviser produces automatically the buy and sell trades necessary to implement the financial advice. An investor can place the trades automatically in batch mode by simply clicking the option on the screen. This feature contributes to making the optimizer highly accessible even to less financially and tech-savvy investors.

An interesting feature of the portfolio optimizer is that it performs an "educational" purpose that can be viewed as an intervention improving financial literacy about risk and return. This is achieved through a data visualization tool that depicts the efficient frontier for the investor. The tool shows the investor both the position of the current portfolio and the position of the proposed portfolio if the optimizer were used. Investors can opt to use the set of stocks held by the investor plus up to 15 additional stocks that represent (in the brokerage firm's view) the most liquid stocks in the Indian stock market each day. Diversification can be achieved by modifying the existing weights of the portfolios and by increasing the number of stocks.

The robo-adviser we analyze is similar to the Portfolio Visualizer marketed in the US by Silicon Cloud Technologies,[8] and is specifically catered to investors that are interested in selecting individual securities, rather than holding ETFs. It displays differences with respect to popular robo-advisers marketed in the US. First, it uses only individual stocks. Second, although it imposes short-sale constraints and operates shrinkage on the estimated variances and co-variances, it uses only 3 years of data for estimation. US-based robo-advising companies do not report the horizon of the data they use, but the 3 years might deliver unstable covariance estimates. The optimizer is also likely to overweigh momentum stocks that have appreciated substantially over the previous years in the proposed portfolio. Finally, no strict rule exists to identify the 15 additional stocks the optimizer considers to add to the investor's portfolio upon usage.

We would like to stress that our analysis does not aim to provide the optimal design for a robo-adviser Rather, our aim is to estimate the causal effects of a robo-advising tool that is available to investors in the field – and whose main features make it similar to its U.S. counterparts – on performance and decision-making.

4 Data

We use four main datasets. Table 1 reports baseline demographic information (age, gender, and account age) for our full sample, as well as for the subsamples we use in the analysis – as described below.

The *Portfolio Optimizer dataset* collects all the individual instances in which a client of the brokerage house used the portfolio optimizer, from the date in which the optimizer was first introduced as an option to clients, that is, July 14, 2015, until February 17, 2017. For each instance, we observe the unique client identifier, the date and time of usage, and the ticker identifier and weight for each of the stocks in the optimizing portfolio. Figure 1 plots the overall number of portfolio optimizer requests each week (dashed line, left y-axis), as well as the first-time requests by each investor (dashed line, right y-axis). Requests peaked in July 2015, when the tool was introduced for the first time and heavily marketed to clients, and in July 2016, once the brokerage house ran a massive round of advertising and marketing of the tool to their clients. On top of these company-wide promotion campaigns, the

---

[8]For further information, see `https://www.portfoliovisualizer.com`

company asked each day different advisers to contact their clients and promote the use of the portfolio optimizer. The average weekly number of requests was around 2,000 over the period, of which about 1,200 were first-time requests.

The second dataset we use – *Transactions dataset* – collects the full trading history of each client of the brokerage house from April 1, 2015 until January 27, 2017. In this dataset, we observe the unique client identifier, the date and time of any transaction made by the client, the ticker of the company on which the client traded, the type of trade, the rupee amount and quantity of the stock traded, the market price of the stock at the time of the trade, whether the trade was executed through the adviser or autonomously by the investor, and the fees charged to the investor. Matching the *Transaction dataset* to the *Portfolio Optimizer dataset* allows us to study the trading behavior of each investor before and after the adoption of the portfolio optimizer.

The third dataset we use – *Holdings dataset* – collects the monthly asset holdings for each client. For the holdings, we observe the unique client identifier, the exact date and time at which the holdings snapshot was registered, the ticker of each security held, the quantity of the security held, and the overall number of assets in the portfolio. The *Holdings dataset* is only available from January 1, 2016 to January 1, 2017.

The last dataset we use – *Logins dataset* – includes all the instances in which an investor or the investor's human adviser connected to the investor's online account. For each login, we observe the date and time in which the account was accessed, whether the investor himself or his/her advisor accessed the account, and whether the access was successful or not. The login information is available for the period between April 1, 2015 and January 27, 2017.

5   Selection into the Adoption of Robo-advising

In the first part of our analysis, we study the selection of individual investors into adopting the robo-advising technology. We perform a simple cross-sectional comparison across two groups of clients of the brokerage house, that is, users and non-users. We start from the raw data, and we restrict the analysis to the sample of investors that place at least one trade during our sample period. We compare the demographic characteristics and trading performance of investors that adopt and do not adopt the robo-advising tool at any point in time since July 2015 – when the brokerage house first introduced

the tool.

Panel A of Table 2 compares the time-invariant characteristics of investors that adopt the robo-advising tool to those that do not adopt the tool, whose trading activity we observe over the same period. Adopters are slightly older than non-adopters, but we cannot reject the null that there is no difference. The average age of adopters is 46.2 years (median: 44.9 years), whereas the average age of non-adopters is 47.8 years (median: 46.9 years). The two groups are similar with respect to the other time-invariant characteristics we observe. The average fraction of men is 71% in both samples, and the average age of the account is 5.8 years in both sample. Overall, we fail to detect any economically or statistically significant difference in time-invariant demographics between users and non-users.

Table 2 also reports the main outcome variables across adopters and non-adopters. Panel B focuses on investors' attention and trading behavior. Portfolio optimizer users are more attentive to their accounts. They login to their online accounts on average 658 times throughout our sample period, whereas non-users log in on average 433 times. Users also place more trades on average (186 vs. 122), have a higher volume of trades (10.6 million rupees vs. 6.0 million rupees), and hence produce a larger amount of trading fees (17.7 thousand rupees vs. 10.07 thousand rupees). Overall, users of the robo-advising tool are more active investors.

In Panel C of Table 2, we compare the trading performance of users and non-users, whereas in Panel D we compare the characteristics of their portfolios at a specified date – January 1, 2016. Two patterns emerge. First, users have a substantially higher amount of assets under management (AUM) and hold more stocks than non-users. Differences are still detected but less substantial when comparing AUM and number of assets for non-stock securities, such as bonds, mutual funds, and ETFs. These other securities represent mere fractions of the value of the stock portfolios investors hold in our sample. Second, Panel C suggests users outperform non-users over our sample period, although both underperform with respect to the market. The 1-month market-adjusted returns of stocks purchased are on average -0.86% for users and -1.22% for non-users. The 3-month market-adjusted returns are on average -2.55% for users and -3.60% for non-users.

The better trading performance of users despite their higher trading activity suggests that users might be more experienced and savvy than non-users. To assess this conjecture in the raw data, we compare the ex-post performance of the stocks purchased to the ex-post performance of the stocks sold. This comparison is based on Odean (1999), who document that the stocks individual investors

sell tend to outperform the stocks they buy. As a rough measure of performance, we compare the market-adjusted returns at 1 and 3 months for the stocks each group of investors purchases and sells. As conjectured, users of the robo-advising tool seem less prone to sell future outperformers than non-users. The difference between the returns of stocks sold minus bought at the 1-month horizon is 0.44 percentage points for users, and 0.55 percentage points for non-users. The difference at the 3-month horizon is 0.76 percentage points for users, and 1.06 percentage points for non-users.

Overall, users of the robo-advising tool do not differ substantially from non-users in terms of demographic and time-invariant characteristics, but they appear to be more sophisticated and to have a higher amount of AUM as well as higher trading activity than non-users.

6 Robo-advising, Trading Behavior, and Performance

In the second part of the analysis, we study the effects of using the robo-advising tool on investors' holdings, trading behavior and trading performance. In this section, we restrict the sample to investors that use the portfolio optimizer at any point in time since July 2015. For those that use the optimizer more than once, we only consider the first date of usage of the optimizer.[9] Our baseline design is a single-difference approach, in which we compare investors' trading behavior and performance before and after the first usage of the optimizer. This single-difference approach allows us to ensure that no time-invariant characteristics of investors explain any changes in trading behavior and performance.

6.1 Robo-advising and Portfolio Diversification

The first set of outcomes we consider are diversification outcomes, that is, the number of stocks investors hold as well as the volatility of their portfolios.

We argue that three – not necessarily mutually exclusive – hypotheses predict robo-advising might increase the portfolio diversification of investors even if human advisers were not able to achieve this goal. First, human advisers might be unaware of the concept of diversification, and might themselves display behavioral biases that they transfer to their clients while advising them. This interpretation is in line with the results of Linnainmaa, Melzer, and Previtero (2017), and could be consistent with our results. Second, financial advisers may encourage their clients to diversify their portfolio, but doing

---

[9]The median user of the portfolio optimizer uses it once.

so requires approving multiple trades in multiple stocks. The process of approving multiple trades is complex, because clients are likely to have concerns regarding the quality of each stock traded and the soundness of each trade. The discussion can become lengthy and quickly lead to paralysis. Instead, following the advice of the robo-adviser is easy and simple. The customer only has to click on a button after seeing the proposed portfolio allocation. Third, human advisers might be aware of the advantages of diversification and that investors display behavioral biases, but they might decide to cater their advice to the tastes of their clients instead of correcting their mistakes and helping them improve their performance.[10]

Table 3 reports the average change in a set of portfolio-level outcomes across investors before and after usage. Panel A reports the average change in the number of stocks (column (1)) and in the market-adjusted portfolio volatility (column (2)). In column (1), we find that on average investors increase the number of stocks they hold by 0.16 units, which is about 1.3% of the median number of stocks investors held before using the portfolio optimizer (12 stocks).

Pooling together all investors masks substantial variation of the effects in the cross-section, especially based on the extent of diversification before using the optimizer. Table 2 shows that some investors are underdiversified – e.g., they only hold 1 or 2 stocks – whereas other investors hold a large number of stocks. For investors that are diversified and hold a large number of stocks to begin with, the optimizer should not necessarily recommend an increase in the number of stocks. If anything, the optimizer might set some optimal weights to zero because of short-sale constraints. Based on this conjecture, we would expect that the number of stocks held increases for underdiversified investors after using the optimizer, and portfolio volatility decreases for them, whereas both dimensions do not change for investors that were diversified before using the optimizer.

To assess the effect of the portfolio optimizer on diversification conditional on the extent of diversification before usage, we first compute the difference between the number of stocks each investor holds in the month after the first usage of the portfolio optimizer and the average number of stocks they held in the month before the first usage of the portfolio optimizer. We then compute the average difference separately for 4 groups of investors, based on the number of stocks they held before usage.

The top panel of Figure 2 reports the results of this exercise. Bars represent the average difference

---

[10]Our setting does not allow disentangling the role of these potential explanations, which we believe future research should investigate.

between the number of stocks held after and before the first usage of the optimizer, which is measured on the y-axis. On the x-axis, we sort investors in 4 groups based on the number of stocks they held before using the optimizer. We report 90% confidence intervals around the estimated means.

Consistent with the conjecture described above, the association between the pre-usage number of stocks and the change in the number of stocks held after usage displays an evident monotonic pattern. Investors that held 1 or 2 stocks before using the optimizer, and hence had the largest need to diversify their portfolio, increase the number of stocks they hold substantially after the first usage of the optimizer. This group of investors increases the size of the portfolio by about 100% on average. The effect is positive both economically and statistically also for those holding between 3 and 5 stocks and between 6 and 10 stocks, but the estimated magnitudes of the change decrease significantly the higher the number of stocks held. Finally, the change becomes negative and statistically significant for those holding more than 10 stocks, which is consistent with the notion that the optimizer might suggest to disinvest from stocks that should be shorted absent the short-selling constraint.

We move on to assess the effects of using the portfolio optimizer on the market-adjusted risk of investors' portfolios. Market adjusted risk is the difference between portfolio realized volatility and market realized volatility at the monthly level, both computed using daily data. In column (2) of Table 3, we consider all investors. We find that on average market-adjusted portfolio volatility decreases by 2.07% per year.[11]

Again, the average result masks substantial heterogeneity based on the ex-ante levels of diversification. In the bottom panel of Figure 2, each bar represents the change in the market-adjusted risk of investors' portfolios across our 4 groups sorted on the number of stocks held before using the optimizer. Consistent with the results on the change in the number of stocks held, we uncover a monotonic pattern whereby abnormal portfolio volatility decreases substantially for investors that held 1 or 2 stocks before using the optimizer. The extent of the decrease in volatility is significantly lower for investors that held between 3 and 5 stocks, and it is even lower for investors that held more than 5 stocks. Note that whereas investors that were diversified ex ante decrease the number of stocks held, their market-adjusted risk does not increase, which suggests that the portfolio optimizer increases portfolio diversification also for this group.

To further assess the extent to which adopting the robo-advising tool affected investors' holdings,

---

[11]We annualize the coefficient in column (2) multiplying it by $\sqrt{12}$.

we compute the share of investors that changed their portfolio holdings within each group – extensive margin.

In figure 3, the left y-axis measures the share of investors that increase the number of stocks they hold after adoption compared to before, for each of the 4 groups sorted by the number of stocks investors held before adoption. This axis is associated with the solid, black line. The right y-axis measures the share of investors that decrease the number of stocks they hold after adoption compared to before. The right y-axis is associated with the dashed, blue line.

Figure 3 shows that the extensive margins of the increase and decrease of stock holdings after adoption of the robo-advising tool are in line with the intensive-margin analysis described above. On the one hand, the share of investors that increase their stock holdings after the adoption of the robo-advising tool is about 38% among the investors that held less than 3 stocks before adoption. This share decreases monotonically the higher the number of shares held before adoption, and is about 22% for investors that held more than 10 stocks before adoption. On the other hand, the share of investors that decrease the number of stocks held after adoption grows from about 5% for the least diversified to 24% for those that held more than 10 stocks before adoption.

Overall, the within-investor single-difference analysis suggests that the portfolio optimizer does increase portfolio diversification for those investors that need diversification at the time they use the tool. Instead, the optimizer does not change the number of stocks held – or, if anything, it decreases it – for those investors that hold more than 10 stocks. Consistently, market-adjusted portfolio volatility decreases substantially for ex-ante less diversified investors, and this decrease declines monotonically with the number of stocks investors held before using the optimizer.

6.2   Robo-advising, Investment Performance, and Trading Activity

We move on to consider investment performance and trading activity. As far as investment performance is concerned, we consider both market-adjusted portfolio performance and the market-adjusted returns of individual trades (stock purchases). For trading activity we consider the overall amount of brokerage fees investors pay, which captures trading volume, and the amount of attention investors allocate to their portfolios, as proxied by the number of days with logins to their online brokerage accounts.

17

Panel B of Table 3 reports the average change in investors' market-adjusted trade performance (column (1)) and market-adjusted portfolio performance (column (2)). In both cases, the average change is positive, although we can reject the null that the coefficient equals zero at plausible levels of significance only for the market-adjusted portfolio performance.

Figure 4 shows the estimation separately across groups of investors, based on the number of stocks they held before using the optimizer. Both average trade performance (top panel) and average portfolio performance (bottom panel) improve significantly for the investors that were underdiversified before usage. At the same time, performance does not change significantly, either economically or statistically, for the other groups.

As far as trading activity is concerned, in Panel C of Table 3 we report the average change across all investors of our proxy for trade volume (column (1)) and the overall number of days with logins (column (2)). On average, monthly fees increase by 155 rupees, which is about 15% of the average amount of fees investors paid in the month before using the optimizer (1,000 rupees). Moreover, on average users of the portfolio optimizer login to their online account 10 days in the month before adoption, and we find that on average they increase this figure by 1 day, which is 10% of the average effect.

When we split the set of investors based on ex-ante diversification, we find again substantial heterogeneity across groups. The top panel of Figure 5 shows that trading fees only increase significantly, both economically and statistically, for investors that were already diversified before using the portfolio optimizer.

When we split the effect of using the optimizer on the number of days with logins, we find that all investors pay more attention to their portfolios, irrespective of the number of stocks held before using the optimizer (see the bottom panel of Figure 5).

6.3   Robo-advising and Behavioral Biases

The last set of outcomes we study relates to a set of well-documented biases attributed to individual investors by earlier research. On the one hand, the trades suggested by the robo-advising tool should not reflect any behavioral biases.[12] On the other hand, because investors trade more after using the

---

[12]Note that recent research suggests human advisors might themselves be subject to behavioral biases, and hence transmit such biases to the trading behavior of their clients (see Linnainmaa, Melzer, and Previtero, 2017) Because

robo-advising tool, the effects of behavioral biases could be higher if investors increased the number of trades they placed without a direct recommendation by the robo-adviser.

We focus on three well-documented behavioral biases, that is, (i) the disposition effect, whereby investors are more likely to realize gains than losses on their positions; (ii) trend chasing, whereby investors tend to purchase stocks after a set of positive returns with the expectations that positive returns will be more likely than negative returns going forward; and (iii) the rank effect, whereby investors are more likely to sell the best performing and worst performing stocks in their portfolios, compared to the other stocks. We find the three biases are substantially less pronounced for all investors after usage of the portfolio optimizer, irrespective of their level of diversification before usage. At the same time, the tool does not fully debias investors.

6.3.1 Disposition Effect (Gambler's Fallacy)

The disposition effect is the tendency to realize gains more often than losses (e.g., Odean (1998)). To measure the extent of disposition effect in our sample, we compute the difference between the proportion of gains realized (PGR) and the proportion of losses realized (PLR) for all investor-days before and after using the portfolio optimizer,[13] where:

$$PGR = \frac{Realized\ Gains}{Realized\ Gains + Paper\ Gains}$$

$$PLR = \frac{Realized\ Losses}{Realized\ Losses + Paper\ Losses}.$$

Investors display a disposition effect if PGR>PLR. Moreover, the larger the positive difference between PGR and PLR, the more severe the disposition effect the investor displays. The disposition effect is an example of *gambler's fallacy*: investors sell gaining stocks because they expect gaining stocks to lose going forward; at the same time, investors do not want to sell losing stocks because they expect them to rebound and gain more going forward.

In the top left panel of Figure 6, each bar represents the difference between PGR and PLR as

---

robo-advising algorithms are designed by humans, these algorithms might themselves reflect the behavioral biases of those designing them.

[13]Note that Odean (1998) only uses days with trades in the computations.

defined above. The bar to the left is the average difference across portfolio optimizer users before usage, whereas the bar to the right is the average difference after usage. Two results are apparent from the figure. First, the extent of disposition effect decreases after usage of the robo-advising tool At the same time, the difference between PGR and PLR is still statistically and economically greater than zero even after usage, suggesting robo-advising does not fully debias investors.

The average difference PGR-PLR is significantly lower after using the portfolio optimizer compared to before using it. We also report a formal test for whether the difference PGR-PLR changes systematically before and after the use of the portfolio optimizer. In Table 4 panel A, we do reject the null that the difference equals zero both statistically and economically at conventional levels of significance.

To assess the economic magnitude of this change, we compare the size of the change to the extent of the bias the average investor displays before using the portfolio optimizer. The size of the difference between PGR and PLR before using the optimizer is about 2 percentage points in our sample. The size of the change of this difference after using the optimizer in Table 4 panel A is about 0.6 percentage points. Therefore, using the portfolio optimizer is associated with a drop in the extent of the disposition effect, as measured by the difference between the proportion of gains realized and the proportion of losses realized, by about 30%, which appears to be a significant economic magnitude.

The limited time span of our data does not allow us to test whether the effect of using the portfolio optimizer on the extent of investors' behavioral biases increases over time – for instance, because investors learn about their biases after understanding they should realize losses if needed – or whether this effect is a one-time shock to the extent of biases. If learning played any role in explaining our results, we might expect that over time the extent of detected behavioral biases would decrease even more than what our current results suggest.

To analyze whether the effects of the optimizer vary across users with different characteristics, we provide two extensions of the baseline results. First, we limit the analysis to those clients that display a positive disposition effect before using the optimizer. We sort these clients into four quartiles, from low incidence of bias to high incidence of bias. We then compute the percentage of clients for which the disposition effect improves – that it, its incidence *decreases* – after using the optimizer. Panel A of Figure 7 reports these percentages in the form of bars. Moving from left to right, the percentage of clients for which the disposition effect decreases moves from approximately 60% for those with a

low incidence of the disposition effect to almost 85% for those with a high incidence of the disposition effect. The confidence intervals attached to each estimated percentage indicate that the effect increases monotonically.

In the second extension we also limit the analysis to those clients that display a positive disposition effect before using the optimizer, but we sort them into four groups based of the number of stocks investors hold at the time of usage. The results reported in Panel B of Figure 7 show that the percentage of clients that experience a decrease in the incidence of their disposition effect is concentrated among the clients with up to five stocks: the percentages are 60% and 65% among the clients with 1-2 and 3-5 stocks, respectively. The effects are less pronounced among customers with more than five stocks.

6.3.2 Trend Chasing

As a second example of investor bias, we consider trend chasing, that is, investors' tendency to purchase stocks after a set of subsequent increases in price, which suggests that investors believe a stock's price is more likely to increase than to decrease after a set of increases.

To measure the extent of trend chasing by investors in our sample, we limit the sample to stocks investors purchase. For each purchased stock, we consider the 5 business days before the purchase date, and we compute the number of days with positive stock returns in this pre-purchase period. We compute the same metric for all stock purchases after investors use the portfolio optimizer:

$$Trend\ Chasing = \frac{Days\ Price\ Increase}{Days\ Price\ Increase + Days\ Price\ Decrease}.$$

The top right panel of Figure 6 plots the average number of days with price increases, both before using the portfolio optimizer (left bar) and after using the optimizer (right bar). On average, before using the optimizer, we do not detect any trend chasing. After usage, we detect a reduction in trend chasing among those that display this bias, therefore lowering the unconditional average. We test formally that the difference between the number of days with price increases before and after the use of the portfolio optimizer is significantly negative in Table 4. In panel B, we reject the null that the within-investor difference equals zero at any standard level of significance.

Similar to our procedure for the disposition effect in the previous subsection, we assess the magnitude of the change in trend chasing by comparing it with the average extent of the bias before usage of the portfolio optimizer. The share of positive returns observed in the 5 days prior to purchase is about 2.45, whereas the size of the change after using the optimizer compared to before in Table 4 panel B is about 0.03. The extent of reduction in trend chasing is thus about 1.2%. The size of this effect appears to be substantially smaller than the effect of using the optimizer on the measure of the disposition effect we proposed in the previous subsection.

6.3.3 Rank Effect

The third bias we consider is the rank effect first documented in a sample of US investors by Hartzmark (2014). The rank effect is the tendency of investors to sell the best and the worst performing stocks in their portfolios, while ignoring stocks in their portfolios that display intermediate performance. To compute the extent of rank effect at the investor level, we follow Hartzmark (2014) and first compute the proportion of best-, worst-, and middle-performing stocks investors sell:

$$Best = \frac{Best\ Sold}{Best\ Sold + Best\ not\ Sold}$$

$$Worst = \frac{Worst\ Sold}{Worst\ Sold + Worst\ not\ Sold}$$

$$Middle = \frac{Middle\ Sold}{Middle\ Sold + Middle\ not\ Sold}.$$

For each investor, we then compute the differences $Best - Middle$ and $Worst - Middle$, both before and after usage of the portfolio optimizer. Under the rank effect, we expect that both differences are statistically different from zero.

In the bottom left panel of Figure 6, each bar refers to the average difference $Best - Middle$ in our sample, both before usage of the optimizer (left bar) and after usage of the optimizer (right bar). In the bottom right panel of Figure 6, each bar refers to the difference $Worst - Middle$, again both before (left) and after (right) usage of the optimizer. As far as the tendency to sell the best performing stocks more than other stocks is concerned, we find that this tendency is substantially higher before

usage than after usage. Similar to the results on the disposition effect, although the extent of the bias decreases after usage of the optimizer, it does not completely fade.

Different from the tendency to sell the best performing stocks, we find that the tendency to sell the worst performing stocks more than the mid-performing stocks is quite limited in our sample. Because of this small baseline effect, we fail to detect any systematic differences in the sizes of the effects for the average investor before and after using the portfolio optimizer.

We confirm these results by testing formally that the change in the prevalence of the rank effect before and after the use of the portfolio optimizer is significantly negative in Table 4. In panel C, we reject the null that the within-investor difference of $Best\ Middle$ equals zero at the 1% levels of significance. To the contrary, in panel D we fail to reject the null that the within-investor difference of $Worst - Middle$ equals zero at any plausible level of significance.

In terms of magnitude of the effect, we note that the share of best performing stocks sold on average before using the optimizer is 22%, whereas the change in this share after using the optimizer compared to before (Table 4 panel C) is 5.7 percentage points, which amounts to about 26% of the average extent of bias. Similar to the disposition effect – of which the rank effect can be considered a special case – the extent to which the portfolio optimizer reduced the bias appears to be substantial.

Overall, the results on behavioral biases suggest that usage of the portfolio optimizer reduces the prevalence of well-known biases among individual investors, although these biases do not wash away completely after the robo-advising intervention.

7   Identification Strategy

All the results we have described so far are based on single-difference tests. This empirical design does not allow us to address a set of alternative explanations for our results. In particular, unobserved time-varying shocks to investors' trading motives might cause both adoption and change in trading behavior.

For instance, an investor could decide she wants to trade more, and might think using the portfolio optimizer will give him/her ideas on which trades to place and how much to invest. If the investor were not to change her behavior without using the optimizer, the single-difference tests would still

estimate the causal effect of the use of the optimizer on investment behavior. But if the investor were to change her trading behavior, had the optimizer being available or not, our baseline results would be spurious.

To address these concerns, we propose an identification strategy that exploits the quasi-random variation of the likelihood that otherwise similar investors use the portfolio optimizer at the same point in time. We build on the fact that the brokerage house asked human advisers to call their clients to promote the usage of the portfolio optimizer and help them use the tool for the first time. The brokerage house had no underlying motivations for pushing the usage of the portfolio optimizer at any point in time, apart from the fact that their technology team thought the device was ready to use broadly and they wanted to market it as an free perk to their clients.

Crucially for our purposes, we observe all the outbound and inbound calls human advisers have with clients at each point in time. Moreover, we know whether calls went through and, if yes, the length of each call. We can therefore construct a treated and a control sample of clients as follows. Treated clients are those clients human advisers reached in the days in which they promoted the portfolio optimizer, and used the optimizer during the call. Control clients are those clients human advisers *tried* to contact on the same days, but did not answer the phone, and hence did not have the chance to hear the adviser promote the tool.[14]

We propose the following difference-in-differences design:

$$(\overline{Outcome}_{reached_t,\ post} - \overline{Outcome}_{reached_t,\ pre}) - (\overline{Outcome}_{missed_t,\ post}) - \overline{Outcome}_{missed_t,\ pre}), \quad (1)$$

where *Outcome* is each of the measures of portfolio diversification, trading performance, and trading activity we studied in Section 6; $reached_t$ indicates investors that were reached by the human adviser on day t; $missed_t$ indicates investors that the human adviser tried to reach on the same day t, but for which the call did not go through; *pre* and *post* refer to the average of each outcomes for the observed period before and after day t.

Our identification strategy translates into the null hypothesis that the quantity defined in (1) equals zero. The crucial identifying assumption is that, absent the usage of the portfolio optimizer,

---

[14]We require that non-responsive clients did not use the portfolio optimizer in the thirty days after the attempted call by their human adviser. The results are not sensitive to using different horizons for this restriction.

the trading activity and performance of investors that were reached on day t would have followed parallel trends with respect to the trading activity and performance of investors that were missed on day t. Under this assumption, missed investors represent a viable counterfactual for the trading activity and performance of contacted investors that used the portfolio optimizer on day t.

Note that the set of clients advisers called may not be random. Advisers might call clients whose characteristics make them more likely to adopt the optimizer, or clients they think would benefit the most from using the optimizer. But this potential selection is not a relevant concern for our strategy, and does not represent a threat to identification in this context. Under such selection, the clients the adviser missed on the portfolio-optimizer promotion days were chosen based on their likelihood of using and/or benefiting from the optimizer, exactly as the clients the adviser was able to reach. Therefore, this type of selection would – if anything – help the econometrician as it would make the treated and control samples similar based on potential unobservable dimensions that determine their trading activity and performance, which the adviser can observe but the econometrician cannot observe.

Another potential concern with our strategy is that it might estimate the causal effect of human advisers suggesting clients they should change their investment strategies, as opposed to the effect of the robo-advising on clients' investment behavior and performance. This concern is barely relevant to our case, because advisers contact clients frequently with their own advice regarding clients' strategies even in days in which they are not promoting the portfolio optimizer. If investors changed their behavior to follow human advisers' recommendations they would have changed their behavior in earlier occasions in which they interacted with their human advisers, and hence we should detect no effects of using of the optimizer.

We estimate the following linear equation by OLS:

$$Change\ Outcome_i = \alpha + \beta \times Treated_i + \epsilon_i, \qquad (2)$$

where $Change\ Outcome_i$ is the difference of the average of each outcome we consider before and after the day in which the adviser tried to reach client i, and $Treated_i$ is an indicator for whether the advisor was able to reach investor i via phone and the investor used the portfolio optimizer that day.

In Table 5, we report the estimated coefficients $\hat{\beta}$ for the set of outcomes we discussed in section 6.

Note that the size of the identification sample is smaller than the size of the baseline sample, because we consider only investors advisers called when promoting the optimizer.

Across all panels, most of the results are qualitatively similar to our baseline results. Two exceptions stand out. First, in column (1) of panel A the coefficient on the change in number of stocks is positive but not statistically different from zero for treated investors. As we saw in the baseline results, though, this coefficient masks dramatic differences in the size of the effect across investors, based on the extent of their ex-ante diversification. We find the same exact monotonic pattern in the identification sample. Figure 8 reports the estimated $\hat{\beta}$ separately for 4 groups of investors, based on the number of stocks they held before using the optimizer. In line with the baseline results, the number of stocks increases significantly for treated investors that held less than 5 stocks before usage. At the same time, the change in the number of stocks is not different from zero for treated investors that held between 6 and 10 stocks, whereas the number of stocks decreases for investors that held more than 10 stocks.

The second departure from the baseline result is the insignificant effect of robo-advice on the portfolio market-adjusted risk, which we report in column (2) of Table 5. In the baseline analysis, we found this effect was negative, whereas in the identification sample we fail to reject the null that the effect is zero, either statistically or economically. Moreover, this effect does not appear to vary systematically across groups of treated investors based on their ex-ante diversification.

Moving on to the identification results for behavioral biases, Table 6 reports the estimated difference-in-differences effects. Our baseline results go through in the identification sample. Specifically, after using the portfolio optimizer, treated investors are less likely to display the disposition effect, less likely to display trend-chasing behavior, and less likely to display the rank effect – although, similar to the baseline analysis, the rank effect is limited to the tendency to sell the best performing stocks.

8   Conclusions

We use a unique sample of individual brokerage accounts to assess the effects of using a robo-advising tool – a portfolio optimizer that makes action on advice simple and immediate – on investor performance and trading behavior, including well-documented behavioral biases.

Adopting robo-advice has different effects across investors based on their extent of diversification

before adoption. Underdiversified investors increase their portfolio diversification in terms of both the number of stocks they hold and the market-adjusted volatility of their portfolio. They display higher performance in terms of market-adjusted trade returns and market-adjusted portfolio returns. Instead, investors that are highly diversified before adoption do not change their diversification. They trade more and their higher trading activity does not translate into better performance.

The extent to which investors are subject to well-known behavioral biases such as the disposition effect, trend chasing, and the rank effect, is the only outcome that improves for all investors, although biases do not disappear.

Overall, our results have implications for the design of robo-advising interventions, which are becoming ubiquitous all over the world. Financial institutions should target underdiversified investors with robo-advising tools, whereas more sophisticated investors and more diversified investors might display lower fee-adjusted performance after using robo-advisers. Future research should dig deeper into the optimal design of robo-advising interventions tailored to the needs of different categories of investors. Despite the promises of robo-advising, its pitfalls do not make it a panacea for individual investors' underperfomance.

References

Aghion, P., and P. Howitt. 1992. A model of growth through creative destruction. *Econometrica* 60:323–51.

Anagol, S., V. Balasubramaniam, and T. Ramadorai. 2017. Endowment effects in the field: Evidence from india's ipo lotteries. *Unpublished Working Paper* .

Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen. 2015. Organizational barriers to technology adoption: Evidence from soccer-ball producers in pakistan. Tech. rep., National Bureau of Economic Research.

Badarinza, C., J. Y. Campbell, and T. Ramadorai. 2016. International comparative household finance. *Annual Review of Economics* 8:111–44.

Barber, B. M., and T. Odean. 2008. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies* 21:785–818.

———. 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The journal of Finance* 55:773–806.

Barberis, N., and R. Thaler. 2003. A survey of behavioral finance. *Handbook of the Economics of Finance* 1:1053–128.

Bhattacharya, U., A. Hackethal, S. Kaesler, B. Loos, and S. Meyer. 2012. Is unbiased financial advice to retail investors sufficient? answers from a large field study. *The Review of Financial Studies* 25:975–1032.

Black, F., and R. B. Litterman. 1991. Asset allocation: combining investor views with market equilibrium. *The Journal of Fixed Income* 1:7–18.

Bold, T., K. C. Kaizzi, J. Svensson, and D. Yanagizawa-Drott. Forthcoming. Lemon technologies and adoption: Measurement, theory and evidence from agricultural markets in uganda. *The Quarterly Journal of Economics* .

Calvet, L. E., J. Y. Campbell, and P. Sodini. 2009. Fight or flight? portfolio rebalancing by individual investors. *The Quarterly journal of economics* 124:301–48.

Campbell, J. Y. 2006. Household finance. *The Journal of Finance* 61:1553–604.

Campbell, J. Y., and L. M. Viceira. 2002. *Strategic asset allocation: portfolio choice for long-term investors.* Oxford University Press, USA.

Chandra, A., A. Finkelstein, A. Sacarny, and C. Syverson. 2016. Health care exceptionalism? performance and allocation in the us health care sector. *The American economic review* 106:2110–44.

Comin, D., and M. Mestieri. 2014. Technology diffusion: Measurement, causes, and consequences. In P. Aghion and S. N. Durlauf, eds., *Handbook of Economic Growth*, vol. 2 of *Handbook of Economic Growth*, 565 – 622. Elsevier.

Conley, T. G., and C. R. Udry. 2010. Learning about a new technology: Pineapple in ghana. *American Economic Review* 100:35–69.

Dietvorst, B., J. Simmons, and C. Massey. Forthcoming. Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science* .

Dupas, P. 2014. Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica* 82:197–228.

Dupas, P., and J. Robinson. 2013. Why don't the poor save more? evidence from health savings experiments. *The American Economic Review* 103:1138–71.

Frame, W. S., and L. J. White. 2004. Empirical studies of financial innovation: lots of talk, little action? *Journal of Economic Literature* 42:116–44.

Frydman, C., S. M. Hartzmark, and D. H. Solomon. Forthcoming. Rolling mental accounts. *Review of Financial Studies* .

Gargano, A., and A. G. Rossi. 2017. Does it pay to pay attention? *Working Paper* .

Gennaioli, N., A. Shleifer, and R. Vishny. 2015. Money doctors. *The Journal of Finance* 70:91–114.

Grinblatt, M., and M. Keloharju. 2001a. How distance, language, and culture influence stockholdings and trades. *The Journal of Finance* 56:1053–73.

———. 2001b. What makes investors trade? *The Journal of Finance* 56:589–616.

Hartzmark, S. M. 2014. The worst, the best, ignoring all the rest: The rank effect and trading behavior. *The Review of Financial Studies* 28:1024–59.

Ledoit, O., and M. Wolf. 2004. Honey, i shrunk the sample covariance matrix. *The Journal of Portfolio Management* 30:110–9.

Linnainmaa, J. T., B. T. Melzer, and A. Previtero. 2017. Retail financial advice: Does one size fit all? *The Journal of Finance* 72:1441–82.

Manuelli, R. E., and A. Seshadri. 2014. Human capital and the wealth of nations. *The American Economic Review* 104:2736–62.

Markowitz, H. 1952. Portfolio selection. *The Journal of Finance* 7:77–91.

Mullainathan, S., M. Noeth, and A. Schoar. 2012. The market for financial advice: An audit study. Tech. rep., National Bureau of Economic Research.

Odean, T. 1998. Are investors reluctant to realize their losses? *The Journal of finance* 53:1775–98.

———. 1999. Do investors trade too much? *American Economic Review* 89:1279–98.

Romer, P. M. 1990. Endogenous technological change. *Journal of political Economy* 98:S71–S102.

Shefrin, H., and M. Statman. 1985. The disposition to sell winners too early and ride losers too long. *The Journal of Finance* 40:777–90.

Skinner, J., and D. Staiger. 2015. Technology diffusion and productivity growth in health care. *Review of Economics and Statistics* 97:951–64.

Thaler, R., and H. Shefrin. 1981. An economic theory of self-control. *The Journal of Political Economy* 89:392–4006.

Tufano, P. 1989. Financial innovation and first-mover advantages. *Journal of Financial Economics* 25:213–40.

Figure 1: Number of Individual Requests to Use the Portfolio Optimizer over Time

This figure plots the overall number of requests to use the portfolio optimizer by all the brokerage house clients (solid line, left y-axis), as well as the requests to use the portfolio optimizer for the first time (dashed lines, right y-axis), for each week between July 1st 2015 – when the tool was first introduced to the clients of the brokerage house – and January 2017.

Figure 2: Portfolio Diversification and Risk Before and After Robo-advising

This figure documents the change in portfolio diversification and risk by investors that use the portfolio optimizer, before and after usage. In both panels, investors are sorted on the x-axis based on the number of stocks they held before using the robo-advising tool. As for the y-axes, in the top panel we report the change in the number of stocks investors hold in their portfolios one month after usage compared to one month before usage. In the bottom panel, we report for the same groups the change in the market adjusted risk of the investors' portfolio. Market adjusted risk is the difference between portfolio realized volatility and market realized volatility at the monthly level, both computed using daily data. Bars refer to the point estimate of the average values within each category of investors. The vertical segments represent 90% confidence intervals for the true mean values within each category of investors.

Figure 3: Investors that Increase and Decrease the Number of Stocks Held After Robo-advising

This figure documents the extensive-margin changes in the number of stocks held after usage of the robo-advising tool. The x-axis sorts investors based on the number of stocks they held before using the robo-advising tool. The left y-axis is associated with the solid, black line. It reports the fraction of investors within each group, who increased the number of stocks held over the month after the first usage of the robo-advising tool, compared to the month before usage. The right y-axis is associated with the dashed, blue line. It reports the fraction of investors within each group, who decreased the number of stocks held over the month after the first usage of the robo-advising tool, compared to the month before usage.

Figure 4: Investment Performance Before and After Robo-advising

This figure documents the change in investment performance at the individual trades (stock purchases) and portfolio levels by investors that use the portfolio optimizer, before and after usage. In both panels, investors are sorted on the x-axis based on the number of stocks they held before using the robo-advising tool. As for the y-axes, in the top panel we report the change in the three-month risk-adjusted performance of the stock purchases placed in the month after usage, compared to those placed in the month before usage. In the bottom panel, we report for the same groups the change in the market adjusted returns of the investors' portfolio. Market adjusted return is the difference between the investor portfolio return and the market return computed over one month. Bars refer to the point estimate of the average values within each category of investors. The vertical segments represent 90% confidence intervals for the true mean values within each category of investors.

Figure 5: Trading Activity and Attention Before and After Robo-advising

This figure documents the change in trading activity and investor attention by investors that use the portfolio optimizer, before and after usage. In both panels, investors are sorted on the x-axis based on the number of stocks they held before using the robo-advising tool. As for the y-axes, in the top panel we report the change in the trading fees paid in the month after usage, compared to the trading fees in the month before usage. In the bottom panel, we report for the same groups the change in the number of days with logins in the month after usage, compared to the trading fees in the month before usage. Bars refer to the point estimate of the average values within each category of investors. The vertical segments represent 90% confidence intervals for the true mean values within each category of investors.
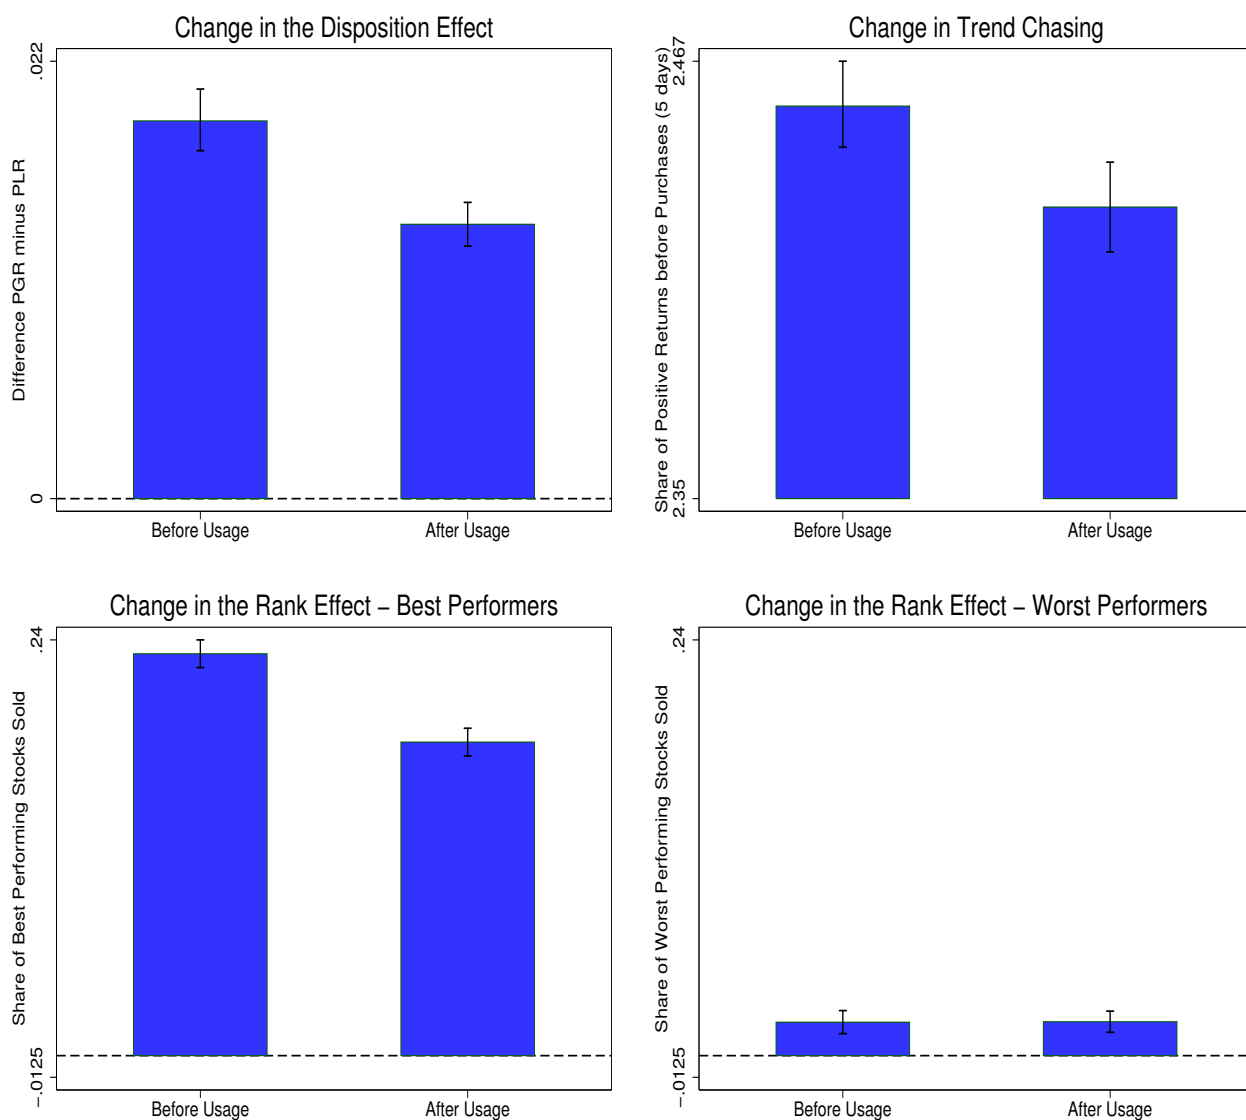
Figure 6: Behavioral Biases Before and After Robo-advising

This figure documents the change in behavioral biases by investors that use the portfolio optimizer, before and after usage. The top left panel reports the results for the disposition effect. Each bar is the average difference between the proportion of gains realized (PGR) and the proportion of losses realized (PLR) for each investor before and after using the optimizer. The top right panel reports the results for trend chasing. Each bar is the average number of days in which a stock purchased by the investor had positive daily returns among the 5 business days before the purchase. The bottom panels report the results for the rank effect for the best performing stocks in the investors' portfolio on the left and the worst performing stocks in the investors' portfolio on the right. Each bar is the average difference between the number of best/worst performing stocks sold and the number of mid-performing stocks sold before and after the use of the optimizer. The vertical segments are 90% confidence intervals for the true mean values within each category of investors.

36

Figure 7: Disposition Effect Before and After Robo-advising – Heterogeneity

This figure documents the change in the disposition effect for investors that use the portfolio optimizer, before and after usage, conditioning on the extent of the disposition effect before usage (Panel A) and on the number of stocks held before accessing the portfolio optimizer for the first time (Panel B). The analysis includes only clients with a positive disposition effect before using the optimizer. In each Panel, clients are sorted into four quartiles, based on the value of the sorting variable. Each Panel reports the percentage of clients that experience an improvement in the disposition effect after usage within each bin. The vertical segments represent 90% confidence intervals around the estimated means.

Figure 8: Identification Results: Number of Stocks and Portfolio Optimizer

This figure reports the identification results for the change in portfolio diversification by investors that use the portfolio optimizer, before and after usage, relative to those investors that were contacted on the same day by the same adviser, but did not answer the phone call. Investors are sorted on the x-axis based on the number of stocks they held before using the robo-advising tool. The y-axis reports the change in the number of stocks investors hold in their portfolios one month after usage compared to one month before usage – relative to the change for the control group. Bars refer to the point estimate of the average values within each category of investors. The vertical segments represent 90% confidence intervals for the true mean values within each category of investors.

Table 1. Demographic Characteristics

A. All Accounts								
	Obs	**Mean**	**St.Dev**	**p.1**	**p.25**	**p.50**	**p.75**	**p.99**
Age	860,943	47.30	13.63	20.73	36.72	45.80	56.80	82.17
Male	838,364	0.75	0.44	0.00	0.00	1.00	1.00	1.00
Account Age	880,254	7.41	3.68	0.12	5.16	8.44	10.12	13.21
B. Accounts with at Least One Trade								
	Obs	**Mean**	**St.Dev**	**p.1**	**p.25**	**p.50**	**p.75**	**p.99**
Age	265,538	46.26	14.14	19.21	35.12	45.02	56.53	80.60
Male	258,656	0.71	0.46	0.00	0.00	1.00	1.00	1.00
Account Age	265,310	5.83	3.96	0.21	1.94	6.08	9.27	13.08
C. Accounts with Holdings Information								
	Obs	**Mean**	**St.Dev**	**p.1**	**p.25**	**p.50**	**p.75**	**p.99**
Age	282,795	48.28	13.32	21.79	38.01	47.28	57.73	81.15
Male	274,048	0.72	0.45	0.00	0.00	1.00	1.00	1.00
Account Age	283,323	7.64	3.27	1.33	5.53	8.38	10.11	13.10
D. Accounts with Logins Information								
	Obs	**Mean**	**St.Dev**	**p.1**	**p.25**	**p.50**	**p.75**	**p.99**
Age	138,482	41.52	13.30	16.98	31.37	38.84	50.35	76.59
Male	136,330	0.74	0.44	0.00	0.00	1.00	1.00	1.00
Account Age	138,405	4.06	3.75	0.12	0.92	2.29	7.04	12.86
E. Accounts that Use the Portfolio Optimizer								
	Obs	**Mean**	**St.Dev**	**p.1**	**p.25**	**p.50**	**p.75**	**p.99**
Age	12,714	48.00	14.49	17.02	36.54	47.10	59.03	81.14
Male	12,386	0.71	0.45	0.00	0.00	1.00	1.00	1.00
Account Age	12,706	6.01	4.09	0.28	1.88	6.06	9.61	13.08

This table presents summary statistics of the demographic characteristics in our datasets. For each variable in each panel, we report the total number of observations (Obs), the sample mean ($Mean$), the sample standard deviation ($St.Dev$) and the $1st$, $25th$, $50th$, $75th$ and $99th$ percentiles of the distributions. Panel A considers all account holders. Panel B considers only those accounts that have traded once over the period April 2015 – January 2017. Panel C considers only account holders for which we have holdings information over the period January 2016 – January 2017. Panel D considers account holders for which we have logins information over the period April 2015 – January 2017. Finally, Panel E considers account holders that use the portfolio optimizer over the period July 2015 – January 2017.

A. Demographic Characteristics

	Non-Users				Users			
	Obs	Mean	St.Dev	Median	Obs	Mean	St.Dev	Median
Age	254,273	46.19	14.13	44.92	11,265	47.81	14.48	46.87
Male	247,674	0.71	0.46	1	10,982	0.71	0.45	1
Account Age	254,053	5.83	3.95	6.09	11,257	5.81	4.09	5.54

B. Attention and Trading Behavior

	Non-Users				Users			
	Obs	Mean	St.Dev	Median	Obs	Mean	St.Dev	Median
Total Logins	98,771	432.85	844.19	84	7,310	657.87	1,020.29	220
Total Trades	254,281	122.38	339.03	15.00	11,265	186.47	398.57	45
Total Volume (Rupee 000)	254,281	5,992	19,181	323	11,265	10,599	25,979	1,196
Total Fees (Rupee 000)	254,281	10.07	27.43	1.09	11,265	17.69	37.03	3.58

C. Trading performance

	Non-Users				Users			
	Obs	Mean	St.Dev	Median	Obs	Mean	St.Dev	Median
Returns Buys (1m)	205,484	-1.22	5.52	-1.11	10,468	-0.86	4.10	-0.86
Returns Sells (1m)	237,395	-0.67	6.38	-0.96	10,797	-0.42	4.81	-0.71
Returns Buys (3m)	201,413	-3.60	10.33	-3.29	10,378	-2.55	7.61	-2.42
Returns Sells (3m)	232,449	-2.54	11.66	-2.77	10,666	-1.79	8.70	-2.22

D. Holdings as of January 1st 2016

	Non-Users				Users			
	Obs	Mean	St.Dev	Median	Obs	Mean	St.Dev	Median
Total AUM	165,983	434,149	1,210,555	72,476	9,327	1,107,550	2,054,217	313,195
Number of Assets	165,983	9.52	12.48	5	9,327	17.27	16.79	12
AUM Stocks	160,402	411,997	1,157,347	68,317	9,208	1,032,630	1,946,557	284,572
Number of Stocks	160,402	9.30	12.27	5	9,208	16.43	16.35	11
AUM Bonds	19,175	141,315	510,280	2,722	2,099	194,415	639,247	5,813
Number of Bonds	19,175	1.61	1.32	1	2,099	1.84	1.64	1
AUM Funds	30,390	78,726	212026	11,890	2,413	125,968	270,957	31,710
Number of Funds	30,390	1.58	1.33	1	2,413	1.97	1.62	1
AUM ETF	8,522	54,158	104,577	18,502	921	63,073	10,9765	22,801
Number of ETFs	8,522	1.19	0.46	1	921	1.30	0.57	1

This table reports summary statistics of the demographic characteristics (Panel A), attention and trading behavior (Panel B), the trading performance (Panel C) and the portfolio holdings (Panel D) of the brokerage account holders in our datasets. In each panel, the results for those that do not use the portfolio optimizer are reported in columns 2 through 5, while the results for those that use the portfolio optimizer at least once are reported in columns 6 through 9. For each variable in each panel, we report the total number of observations (Obs), the sample mean ($Mean$), the sample standard deviation ($St.Dev$) and the sample median ($Median$). The results in panels A through C are computed over the full sample, while the results in Panel D are computed as of January 1^{st} 2016.

Table 3. Diversification, Attention and Trading Behavior Before
and After Adopting the Portfolio Optimizer – Baseline Results

Panel A. Adoption of the Optimizer and Diversification

	Number of Stocks	Portfolio Market Adjusted Risk
Change after Adoption	0.156**	-0.006***
(*p*-value)	(0.04)	(0.02)
Obs	4,672	3,115

Panel B. Adoption of the Optimizer and Investment Performance

	Performance of Trades	Portfolio Market Adjusted Returns
Change after Adoption	0.003	0.005**
(*p*-value)	(0.47)	(0.02)
Obs	1,192	3,428

Panel C. Adoption of the Optimizer, Trading Activity and Attention

	Trading Fees	Days with Logins
Change after Adoption	155.4***	0.853***
(*p*-value)	(0.00)	(0.00)
Obs	6,594	4,000

This table reports results on investor behavior before and after adopting the portfolio optimizer. Panel A reports the changes in the number of stocks held (first column) and the market adjusted risk of the investor portfolio (second column). Panel B reports the changes in the risk-adjusted performance of the stock purchases (first column) and the market adjusted performance of the investor portfolio (second column). Panel C reports the changes in the trading fees paid to the brokerage house (first column) and the number of days with logins (second column). All panels compare the behavior over the month after the usage and the behavior over the month before the usage. Each panel reports first-difference coefficients, the associated *p*-values and the number of observations.

Table 4. Behavioral Biases Before and After Adopting
the Portfolio Optimizer − Baseline Results

	Panel A. Disposition Effect	Panel B. Trend Chasing Behavior
Change after Adoption	−0.006***	−0.027***
(p-value)	(0.00)	(0.00)
Obs	7,506	6,938

	Panel C. Rank Effect − Best	Panel D. Rank Effect − Worst
Change after Adoption	−0.057***	0.007
(p-value)	(0.00)	(0.123)
Obs	4,264	4,264

This table tests whether the change in behavioral biases by investors that use the portfolio optimizer is different from zero before and after usage. Panel A reports the results for the disposition effect. *Change after Adoption* is the difference between the proportion of gains realized (PGR) and the proportion of losses realized (PLR) for each investor before and after using the optimizer. Panel B reports the results for trend chasing. *Change after Adoption* is the difference between the average number of days in which a stock purchased by the investor had positive daily returns among the 5 business days before the purchase, before and after adoption. Panel C and Panel D report the results for the rank effect. *Change after Adoption* is the average difference between the number of best/worst performing stocks sold and the number of mid-performing stocks sold before and after the use of the optimizer. Each panel reports first-difference coefficients, the associated p-values and the number of observations.

Table 5. Diversification, Attention and Trading Behavior Before and After Adopting the Portfolio Optimizer – Identification Results

Panel A. Adoption of the Optimizer and Diversification

	Number of Stocks	Portfolio Market Adjusted Risk
Treated	0.339	0.002
(p-value)	(0.37)	(0.65)
Obs	720	509

Panel B. Adoption of the Optimizer and Investment Performance

	Performance of Trades	Portfolio Market Adjusted Returns
Treated	0.004	0.031**
(p-value)	(0.68)	(0.01)
Obs	815	542

Panel C. Adoption of the Optimizer, Trading Activity and Attention

	Trading Fees	Days with Logins
Treated	318.9*	1.011***
(p-value)	(0.08)	(0.00)
Obs	1,507	1,086

This table reports the identification results for the change in portfolio diversification, trading activity, and performance by investors that use the portfolio optimizer, before and after usage, relative to those investors that were contacted on the same day by the same adviser, but did not answer the phone call. Treated is the estimated coefficient $\hat{\beta}$ from the following equation, which we estimate by OLS:

$$Change\ Outcome_i = \alpha + \beta \times Treated_i + \epsilon_i$$

Panel A reports results for the number of stocks held (first column) and the market adjusted risk of the investor portfolio (second column). Panel B reports the changes in the risk-adjusted performance of the stock purchases (first column) and the market adjusted performance of the investor portfolio (second column). Panel C reports the changes in the trading fees paid to the brokerage house (first column) and the number of days with logins (second column). All panels compare the behavior over the month after the usage and the behavior over the month before the usage. For each difference-in-differences coefficient we report the associated p-value and the number of observations in the regression.

Table 6. Behavioral Biases Before and After Adopting
the Portfolio Optimizer – Identification Results

	Panel A. Disposition Effect	Panel B. Trend Chasing Behavior
Treated	-0.008^{***}	-0.069^{***}
(p-value)	(0.00)	(0.00)
Obs	2,766	2,752

	Panel C. Rank Effect – Best	Panel D. Rank Effect – Worst
Treated	-0.058^{***}	-0.006
(p-value)	(0.00)	(0.27)
Obs	2,621	2,621

This table reports the identification results for the change in behavioral biases by investors that use the portfolio optimizer, before and after usage, relative to those investors that were contacted on the same day by the same adviser, but did not answer the phone call. Treated is the estimated coefficient $\hat{\beta}$ from the following equation, which we estimate by OLS:

$$Change\ Outcome_i = \alpha + \beta \times Treated_i + \epsilon_i$$

Panel A reports results for the disposition effect, where *Change Outcome* is the difference between the proportion of gains realized (PGR) and the proportion of losses realized (PLR) for each investor before and after using the optimizer. Panel B reports the results for trend chasing. *Change Outcome* is the difference between the average number of days in which a stock purchased by the investor had positive daily returns among the 5 business days before the purchase, before and after adoption. Panel C and Panel D report the results for the rank effect. *Change Outcome* is the average difference between the number of best/worst performing stocks sold and the number of mid-performing stocks sold before and after the use of the optimizer. For each difference-in-differences coefficient we report the associated p-value and the number of observations in the regression.