Human-Robot Interactions in Investment Decisions *

Milo Bianchi† Marie Brière‡

March 2024
forthcoming in *Management Science*

Abstract

We study the introduction of robo-advising on a large set of Employee Saving Plans. Differently from many services that fully automate portfolio decisions, our robo-advisor proposes investment and rebalancing strategies, leaving investors free to follow or ignore them. The resulting human-robot interactions occur both at the time of the subscription and over time, as the robot sends alerts when the investor's portfolio gets too far from the target allocation. We show that the robo-service is associated with an increase in investors' attention and trading activities. Following the robot's alerts, investors change their rebalancing behaviors so as to stay closer to their target allocation, which results in larger portfolio returns. Counterfactual returns induced by automatic rebalancing by the robot would be only slightly higher, suggesting that on average the financial cost of letting investors retain control is not large.

*We are grateful to Jean-Edouard Colliard (Editor) and two anonymous Referees for very insightful feedback, as well as to Francesco D'Acunto, Francesco D'Amuri, Laurent Bach, Paul Beaumont, Jean-François Bonnefon, Lisa Bruggen, Giuliano de Rossi, Andreas Hackethal, Michael Haliassos, Ulrich Hege, Fathi Jerfel, Luigi Guiso, Sophie Moinas, Mehrshad Motahari, Peter Pope, Sébastien Pouget, Sofia Ramos, Michael Reher, Alberto Rossi, Patrick Tissot-Favre, and Wanyi Wang for very useful discussions. We also thank Corinne Laboureix, Catherine Leroy, Fabien Thevenot for insightful explanations on the institutional setting and Sébastien Cadot, Arnaud Delavoet, Marc-Antoine Hubert, Anthony Lapeyre and Mohamed Youssef for help in the data gathering process. We acknowledge funding from the Observatoire de l'epargne européenne; Milo Bianchi acknowledges funding from LTI@Unito and ANR (ANR-17-EURE-0010 grant). The data used to carry out this study come from the processing of record keeping and account keeping of AMUNDI ESR employee and pension savings accounts. These data have been analyzed anonymously for scientific, statistical or historical research purposes. We have no material interest that relates to the research described in the paper.

†Toulouse School of Economics, TSM, and IUF, University of Toulouse Capitole, Toulouse, France. E-mail: milo.bianchi@tse-fr.eu

‡Amundi, Paris Dauphine University, and Université Libre de Bruxelles. E-mail: marie.briere@amundi.com.

1 Introduction

Automated financial advisors, often called robo-advisors, are attracting a growing attention both in academia and in the industry. Robots have low operating costs, which may allow reaching a broader set of investors, and they adopt verifiable procedures, which may limit the extent of biased advice (Bianchi and Brière (2022), D'Acunto and Rossi (2023)). As for many applications of AI in finance and in other domains, a fundamental question is whether these robots tend to complement, or rather to replace, investors' reasoning and actions.[1] The extent to which investors keep an active role in their decisions appears as a fundamental dimension when assessing whether and how robo-advisors can improve investors' choices and promote financial capability.[2]

Defining the optimal degree of automation is however not an obvious task. First, while reliance on algorithms seems particularly delicate in the context of financial services, evidence from other domains suggests that algo-aversion can be partly overcome by letting humans and robots interact. For example, Dietvorst, Simmons and Massey (2018) and Burton, Stein and Jensen (2020) show in experimental settings that subjects are more willing to rely on automated advice when they retain some control over the underlying algorithm. Moreover, these interactions can also be useful for promoting investors' learning on how to manage their portfolios (see e.g. Seru, Shumway and Stoffman (2010) on the dynamics of learning by trading and Loos, Previtero, Scheurle and Hackethal (2020) on spillovers across contracts managed by a robot and self-managed contracts). At the same time, letting the human interfere with the robot may limit the effectiveness of the robot and ultimately be detrimental to investors' performance (see for example Ge, Zheng, Tian and Liao (2021) on peer-to-peer lending and Green and Chen (2019) on judges' decisions).

In this paper, we investigate one particular dimension of these trade-offs.[3] Our first question is whether significant changes in investors' decisions may occur in a setting in which the robot provides investment and rebalancing recommendations while investors retain full control on their portfolios. The question is important given the widespread approach of substituting investors' decision-making, through mechanisms such as auto-

[1]See the recent literature on how human-robot interactions should be taken into account when designing AI systems (e.g. Raghu, Blumer, Corrado, Kleinberg, Obermeyer and Mullainathan (2019), Mozannar and Sontag (2020), Bansal, Nushi, Kamar, Horvitz and Weld (2021)).

[2]See e.g. Siddarth, Acemoglu, Allen, Crawford, Evans, Jordan and Weyl (2021) and Brynjolfsson (2022) for a deeper discussion on AI systems that complement or substitute humans and their far-reaching economic and social effects.

[3]Drawing a full picture is beyond our scope. Our data do not allow to compare the take-up of our robot to that of a counterfactual fully automated robot, nor to explore in details the long-run consequences of retaining control in terms of improved financial capability.

matic enrollment and rebalancing, and the outstanding debate about the effectiveness of this approach.[4]

Second, we analyze the consequences of the behavioral changes associated with the adoption of the robot in terms of risk exposure and portfolio returns. This allows to shed light on human-robot interactions that occur over time, when for example investors experience market shocks or new investment opportunities arise. As we show, these interactions are key to understanding the ultimate effects of the robot on financial outcomes. This also allows to compare the behavioral changes associated with the robo-service to counterfactual outcomes investors would have experienced had they fully delegated their decisions, thereby providing some elements of comparison between the two scenarios.

We investigate these questions by exploiting the introduction of a robo-advising service by a large French asset manager. The distinctive feature of this service is that the robot gives advice to the investors while leaving them free to follow or to ignore the advice. Another important feature is that these advises occur both at the time of the subscription, when the robo proposes the initial portfolio allocation, and over time, as the robot sends email alerts when the current portfolio allocation ends up being too far from the target. This makes it different from the more common robo-advisors (discussed below) that automate portfolio investment and rebalancing, and this makes it particularly useful for focusing on human-robot interactions both at the time of the subscription and over time.

The robo-advisor under study was introduced in a large set of Employee Saving Plans in August 2017. The robot is proposed to employers and, if they accept, employees get a notification on the availability of the service and decide whether or not to subscribe it. The robot starts by eliciting information on the investor's characteristics, builds the investor's profile, and proposes a portfolio allocation. If the investor accepts the proposal, the robot implements the allocation. Over time, if the portfolio allocation deviates substantially from the target, the investor is prompted to connect to the platform and to rebalance her portfolio towards the target. Absent the robot, employees self-manage their portfolios without any dedicated advice. We have access to account level data covering from January 2016 to June 2021, aggregated at the monthly level. Our sample contains investors who have accepted the robo-service as well as individuals who have not been offered the service (i.e., non-exposed), individuals who have declined the offer without initiating the profiling process (i.e., non-takers), and individuals who have initiated the profiling process without eventually subscribing to the service (who we call

[4]See for example Madrian and Shea (2001) and Choukhmane (2019) for contrasting views on the effectiveness of automatic enrollment in retirement plans.

robo-curious).

An important challenge for our empirical analysis is that the choice of taking up the robot is voluntary and as such it could be driven by unobserved investors' characteristics that are also related to our outcomes of interest. In this respect, a particularly interesting feature of our data is that it allows comparing the behavioral changes associated with the robot's take-up to those displayed by robo-curious investors. In this way, we can condition on the (possibly time-varying) characteristics that may induce investors to express interest in the robo-service and compare the effects of the take-up relative to just observing the robot's profiling and recommendation.[5] Moreover, the data allow to construct the recommendations that the robot would have sent to robo-curious investors, both at subscription and over time, had they accepted the service. We can then define several counterfactual scenario concerning in particular the alerts sent by the robot when the investor's portfolio gets too far from the target, thereby focusing on behavioral changes occurring around the reception of alert, rather than around the robot's take-up.

We first show that investors who take-up the robot do not view it as a substitute for their own participation. Instead, we observe an increase in investors' attention to their portfolio, as measured by the amount of time spent on the dedicated website. Robo-takers increase the number of connections on the platform by 0.3 per month, which can be compared to the average of 1. Importantly, this effect persists even beyond the time of the robo-subscription.

This increase in attention is associated with an increase in trading activities, and specifically in rebalancing activities that occur after the take-up, as investors experience market shocks or as they face new investment opportunities. In particular, our analysis highlights the importance of the alerts sent by the robot as an effective tool for inducing investors to rebalance their portfolio and stay closer to the target. The extent to which investors react to alerts is interesting as it highlights how the reliance on the robo-service evolves over time, as investors may be prompted to pay attention to their portfolios even if not used to do so, or they may be advised to rebalance their portfolio in a given direction even if tempted to do otherwise.

We exploit the knowledge of the algorithm governing the alerts and construct counterfactual alerts that robo-curious would have received had they taken the robot. We then show, by comparing robo-takers to robo-curious in a standard diff-in-diff specification, that the actual reception of the alert increases investors' attention and propensity

[5] We consider alternative control groups (non-takers or not-exposed investors) in robustness checks so as to further isolate the effects of the robot's take-up from potentially confounding factors.

to rebalance, and it reduces the distance between current and target equity exposure by 4.6%, relative to an average distance (conditional on being alerted) of 6.2%.

These changes in trading behaviors have significant consequences for investors' risk exposure and portfolio returns. We first show that after subscribing to the robot, investors tend to increase their equity exposure. We observe an increase in the equity share by 3%, as compared to an average equity share of 22%. These changes in equity exposure translate into significant changes in portfolio returns. We show that after subscribing to the robot, investors experience an increase in returns net of fees of about 2.6% per year. These figures are essentially the same when considering the impact of the robot on expected returns, which we estimate based on a standard risk factor model. This suggests that the increase in returns induced by the robot can be attributed to its dynamic exposure to standard risk factors.

We further investigate the determinants of the increase in returns by decomposing it between a static effect induced by a change in portfolio allocation occurring at the time of the robo-subscription and a dynamic effect induced by the way in which investors rebalance their portfolios over time. We show that the a significant part of the increase in returns comes from a change in rebalancing behaviors. In order to isolate the dynamic effect, we identify a subset of robo-curious investors whose portfolio allocation when completing the survey is close to the one suggested by the robot. For these investors, the potential changes associated to the robot's take-up would have been essentially those associated with rebalancing behaviors after subscription. Using these investors as reference point, we show that about half of the increase in returns associated with the robo-service can be explained by the way in which investors rebalance their portfolios over time.

Finally, we investigate the potential financial costs, in terms of foregone returns, of letting investors decide whether or not to follow the robot, as opposed to completely automating their portfolio rebalancing. Comparing the difference between the returns experienced by robo-takers and the counterfactual returns they would have experienced with automatic rebalancing, we observe an average loss in annualized returns of about 0.025%. The difference is tiny, relative for example to the 2% difference in returns between robo-takers and robo-curious we have estimated above. This suggests that, at least in our setting, the financial costs of having investors-in-the-loop, as opposed to implementing an automatic rebalancing, are on average not large.

At the same time, our figures reveal an important heterogeneity across investors, which is a common feature in studies documenting the cost and benefits of letting humans overrule algorithmic decisions.[6] In our setting, investors' demographic characteristics do not appear as a key determinant of this heterogeneity; market conditions appear more important. In particular, we show that investors are significantly less likely to follow the robot's recommendations during the bear market periods of October-December 2018 and February-March 2020.

Many of the above results are derived within a specific group of investors, the robo-takers and the robo-curious, who have shown some interest in the robot's service. While as mentioned (and further discussed below) comparing these two groups has several advantages, the generalisability of our findings to other groups of investors remains an open question. Despite this important caveat, we believe our results are encouraging on the possibility to promote human–robot interactions in the field of personal finance. The effects on increased attention and trading are surprising in light of a large body of literature documenting investors' inattention to rebalancing opportunities, especially in long-term investment plans as the ones under study (see Gomes, Haliassos and Ramadorai (2021) for an extensive review). Importantly, low attention can have detrimental effects on returns (Gargano and Rossi (2018)) and at the same time portfolio inertia tends to be more pronounced for investors with lower wealth and lower financial literacy, who also tend to experience lower returns (Calvet, Campbell and Sodini (2009), Bianchi (2018)). Under this perspective, it is remarkable that the robo-service can at least partly limit these tendencies.

We contribute to the debate on how automation can impact financial services, and more specifically to a growing literature on the effects of robo-advising on portfolio choices. D'Acunto, Prabhala and Rossi (2019) study an interactive portfolio optimizer offered by an Indian brokerage house and show it has beneficial impacts on less diversified investors, as it induces them to hold a larger number of stocks, but not on diversified investors. D'Acunto et al. (2019), however, do not focus on human-robot interactions and on the resulting portfolio dynamics, which is a central feature in our analysis. As we show, the dynamic interactions between the robot and the investors are in our setting key to understanding how the robot impacts investors' rebalancing behaviors and performance.[7]

[6]See for example Kleinberg, Lakkaraju, Leskovec, Ludwig and Mullainathan (2018) on the importance of selection, Angelova, Dobbie and Yang (2022) on performance heterogeneity, Noy and Zhang (2023) on the implications for aggregate inequality.

[7]Differently from D'Acunto et al. (2019), in our setting investors do not pick stocks but choose among a menu of funds, which should minimize issues of under-diversification. Rebalancing behaviors may be an equally important source of (under)performance, especially for less sophisticated investors (Bianchi (2018)).

The focus on human-robots interactions also distinguishes our paper from most other contributions, such as Reher and Sun (2019), Loos et al. (2020), Rossi and Utkus (2020), Reher and Sokolinski (2021), that study automated portfolio managers in which portfolio choices over time are fully delegated to the robot (see D'Acunto and Rossi (2020) and Bianchi and Brière (2022) for overviews). Our paper shows that the robot can significantly affect investors' decisions even while letting them retain control over their portfolios. As we emphasize in the concluding remarks, we view investors' active participation as an important tool to promote learning and financial capability, and hence to assess the long-term consequences of the robo-service.

2 Data

The portfolio choices under study concern a large set of Employee Saving Plans. Each year, as part of their compensation, employees receive a sum of money to be allocated across a set of funds offered by the employer. The employer can offer two types of contracts, which differ in the lock-in period: 5-years (*plan d'épargne entreprise*) or until retirement (*plan d'épargne pour la retraite collectif*). Employees can make extra investment in the plan, withdraw money after the lock-in period (or under exceptional circumstances), and freely rebalance their portfolios over time. An individual can simultaneously hold several contracts from past and current employers.[8]

These plans are managed by a large French asset manager. While traditionally employees received no advice on these portfolio choices, the asset manager has introduced a robo-advisor service in August 2017. If the employer subscribes to the robo-service, its employees are informed via email and they have the option to accept it on one or more of their saving accounts. The cost of the service is borne by the employee, and it depends on the value of her account.

The robot starts by eliciting information on the investor's characteristics, and specifically on her risk-aversion (both through quantitative and qualitative questions), financial knowledge and experience (both objective and self-assessed), age and investment horizon. Based on these questions, the robot builds the investor's profile (say, prudent, dynamic,..) and proposes a portfolio allocation. Importantly, the robot's allocation is built within the funds proposed by the employer; that is, investors have access to exactly the same menu of funds with and without the robot.[9]

[8]In our sample, we observe on average 3 contracts per investor.

[9]The robot is programmed to propose an allocation on the part of the portfolio which is not invested

The investor can visually compare the proposed allocation with her current one both in terms of macro categories (proportion of equity, bonds, money market funds, ...) and of specific funds. If the investor accepts the proposal, the robot implements the allocation.[10] Over time, the robot also sends email alerts if current portfolio allocation ends up being too far from the target allocation.

We take advantage of several sources of (anonymized) data. First, we have obtained detailed information on the investment choices. We observe the menu of funds offered by the employer, the allocation chosen by the employee, new investments, rebalancing, and withdrawals. In addition, building on the information on returns of the various funds, we have constructed the returns and various measures of risk of these portfolios (as detailed below). Third, we have extracted information about investors' activities on the platform, both in terms of trading and in terms of digital footprints (number of connections, duration, pages visited). Fourth, for individuals who take the robot, we can observe the score they are given by the robot, the associated profile and suggested allocation, and the alerts the robot may be sending over time to propose new allocations.[11] We provide more details about those variables as we proceed with our analysis below.

We first had access to our data in November 2018. At that point, around 8,000 companies were offered the robo-service, that corresponds to 646,884 employees (out of over 1,9 millions active employees in those plans). Out of them, 189,918 individuals had expressed interest in the robot and started the procedure to receive the service by formally signing a "counselling agreement" in at least one of their account. Out of them, 175,342 individuals ended up not subscribing to the service and we refer to them as robo-curious, while the remaining 14,635 individuals have subscribed to the robot and we refer to them as robo-takers. This corresponds to 18,164 accounts managed by the robot in 770 different firms.

We have extracted the trading records of all individuals who have taken up the robo-services as of November 2018, together with random samples of 20,000 individuals who are "not-exposed" (i.e. employees of companies which do not have access to the service), 20,000 individuals who are exposed but non-takers and 20,000 individuals who are curious. We restrict to individuals who have completed at least one transaction in one of their

in employer's stock, which may have some specificities (e.g. in terms of matching rule) relative other stocks.

[10]Even if the investor accepts the robot's allocation, she is not committed to it in any way, she can change again the allocation right after having taken up the robot.

[11]We observe the overall score assigned by the robot, not the single answers provided by the investor on risk aversion, financial literacy, and investment horizon.

accounts in our sample period. We have obtained the corresponding historical records starting in January 2016 and followed these individuals up to June 2021, which gives us a panel covering the period January 2016 to June 2021, aggregated at the monthly level.

Our sample is representative of the French population of private sector employees. The firms under study are representative of the French population of private firms, and all employees in these firms have access to the saving plans. The average value of the assets invested in the plan is 34,811 euros, the median is 12,918 euros. These figures are comparable to those one can find in representative surveys.[12] This allows us to include in our analysis also small investors, who tend to be underrepresented in studies focusing on stock market participants (say, from a brokerage house).

In Table 1 (Panel A), we report some descriptive statistics on demographic and portfolio characteristics of our investors in the different samples (takers, curious, non-takers, not-exposed). We aggregate variables at the individual level and, for each investor, we consider the average value of the variable before the introduction of the robo-service, between January 2016 and August 2017. Robo-takers display some differences with the other investors and, most importantly, these differences may also occur along unobservable and possibly time-varying characteristics. We address this important issue in several ways throughout the next analysis.

3 Attention, Trading and Alerts

In this section, we first investigate whether the robot's take-up is associated with a significant change in the level of attention investors pay to their portfolios. We then consider the associated changes in trading behaviors, focusing particularly on portfolio rebalancing. We highlight the distinctive role played by the alerts sent by the robot, which prompt investors to pay attention to their portfolio and remain close to their target allocations. Our analysis is motivated by the above-mentioned literature on investors' inattention and portfolio inertia (extensively reviewed in Gomes et al. (2021)).

In our baseline analysis, we explore the behavioral changes associated with the robot in a series of fixed-effects regressions. Since an individual can hold several contracts in the plan, and decide to take-up the robot in one or more of her contracts, we consider

[12]For example, data on household savings report average financial wealth around 60,000 euros and, for those who have access to employee savings' plans, these plans represent on average around 20% of their financial wealth. Sources: Observatoire de l'Epargne Européenne (http://www.oee.fr/files/faits_saillants_-_2020_t2.pdf) and Autorité des marchés financiers (https://www.amf-france.org/fr/actualites-publications/publications/rapports-etudes-et-analyses/les-actifs-salaries-et-lepargne-salariale).

specifications at the contract level. We estimate the following equation:

$$y_{j,t} = \alpha_j + \beta RoboTreated_{j,t} + X'_{j,t}\gamma + \mu_t + \varepsilon_{j,t}, \tag{1}$$

where α_j and μ_t are contract and time fixed effects. As mentioned, our data are aggregated at the monthly level; hence, unless specified otherwise, time t refers to a given year-month. $RoboTreated_{j,t}$ is a dummy equal to 1 if the investor has taken the robot in contract j at time t, and $X_{j,t}$ is a vector of individual and portfolio characteristics.

Unless specified otherwise, our controls include the average equity share and the average returns over the past 12 months, the account value in the previous month, the value of the yearly variable remuneration, a dummy if the variable remuneration was received in the current month and a dummy if the variable remuneration was received in the past month. We double cluster standard errors by individual and time. Our coefficient of interest β measures how, in a given contract j, the outcome $y_{j,t}$ varies with the adoption of the robot, compared to the changes experienced in the control group. In most of this analysis, our control group is defined by the sample of robo-curious; that is, individuals who have initiated the profiling process with the robot while eventually declining the service. We consider alternative specifications in the robustness section. Summary statistics of the main variables used in the analysis are reported in Table 1 (Panel B).

3.1 Attention

We first consider the level of attention that investors pay to their portfolios. As mentioned, we have extracted the login activities made on the dedicated platform, and we observe the number of connections, the number of web pages visited, the number of minutes spent on the platform.

Since these activities are recorded at the individual level, rather than at the contract level, we modify Equation (1) and define the treatment variable as equal to 1 if the investor has taken up the robot in at least one of her contracts. We also include individual, rather than contract, fixed effects. Moreover, our attention variables are only available on a shorter sample, from January 2016 to November 2018.

We report our results in Table 2. Our key observation is that, after having taken the robot, investors spend more time on the platform. In column 1, we observe an increase of 0.28 connections per month (the average is 1). Similar patterns hold with other measures of attention.[13]

One may question whether the increased attention is associated with the robo-subscription or to other events occurring at the same time. A typical event that increases investors' attention is the reception of the remuneration that needs to be allocated across the various funds in the saving plan. Employees typically receive a communication before the reception and they are asked to choose their allocation in the next month. Indeed, we observe an increase in activities on the platform during the month of reception of the remuneration, and if that corresponds to the month of robo-subscription we may confound the two effects. In column 2, we exclude the month before and the month at which the individual has received the variable remuneration. We see that our estimates are only slightly smaller than those in column 1.

A related concern is whether the effects persist also beyond the window of the subscription to the service. In column 3, we exclude the month of the robo-subscription. As intuitive, the estimated effects are smaller in magnitude than the overall effects in column 1, but still significantly different from zero. That is, robo-takers display larger levels of attention also beyond the time of the subscription and the time of reception of the variable remuneration.

In order to check the parallel trend assumption and uncover possible dynamics of those effects, we consider the following regression:

$$y_{j,t} = \alpha_j + \sum_{s=-6}^{5} \beta_s \mu_{t+s} RoboTreated_{j,t} + X'_{j,t}\gamma + \mu_t + \varepsilon_{j,t}, \tag{2}$$

where μ_{t-s} and μ_{t+s} correspond to months before and after the take-up and the other variables are as in Equation (1). In Figure 1, we report the estimated coefficients $\beta_{-6}, ..., \beta_5$ and the associated 95% confidence intervals. We observe no significant pre-treatment differences. We will repeat the same analysis considering other variables of interest in the next sections, we collect the corresponding results in Figure 1.

Overall, these result show that investors do not take the robot as a substitute for their own attention. Rather, the robot is associated with an increased level of attention, which persists even beyond the time of its subscription. These results also contrast a popular approach to automated portfolio management in which investors are induced to stop paying attention to their portfolio, despite that investors who pay more attention

[13] We observe an increase in the number of minutes spent on the platform by 3.9 per month (the average is 6.2) and an increase in the number of web pages visited per month by 4.9 (the average is 6.5).

may obtain a better financial performance (Gargano and Rossi (2018)).

3.2 Trading Activities

We now consider whether the increased level of attention is associated with an increase in trading activities. We focus on pure rebalancing activities, in which investors change their portfolio composition by moving money across funds without increasing or decreasing their total investment, as these are directly affected by the interaction with the robot, as detailed below. These operations are not subject to fees on the part of the asset manager.

In column 4 of Table 2, we observe that subscribing to the robot in a given contract is associated with 0.09 more allocation changes by month, relative to an average of 0.05. The total sum of rebalancing activities includes those arising from the robot's portfolio recommendation at the time of subscription, those arising from the robot's rebalancing recommendation after the subscription, and those directly implemented by the investor. In column 5, we focus on portfolio rebalancing arising from a robo-recommendation after subscription, and observe a significant increase of 0.04 in these activities (explained in more details below). In column 6, we observe that rebalancing activities not induced by the robot are not significantly affected, which shows that the increased trading activities are driven by the direct interaction between the investor and the robot. In the next analysis, we explore the role of the alerts sent by the robot in explaining these results.

3.3 The Role of Alerts

An important feature of the robo-service is that it sends alerts to investors in case their current allocation is far from the target allocation. In case of alert, the investor receives an email stating that there is discrepancy between the current and the target allocation, due for example to a market shock, and she is suggested to connect to the dedicated website to consult her portfolio. Once the investor is connected, the robot proposes to rebalance the portfolio so as to get back to the target allocation and, if the investor accepts, the required adjustment is implemented by the robot.

We are interested in investigating how investors respond to those alerts for two reasons. First, we check whether the alerts are effective in inducing investors to rebalance their portfolio so as to stay closer to their target allocation. It has been shown that, even when investing in funds and not in individual stocks, less sophisticated investors tend to chase trends and as a result their risk exposure displays larger sensitivity to market fluctuations (Bianchi (2018)). Second, investors' reaction to alerts highlights whether they are willing to rely on the robo-recommendation not only at the time of the subscription but also after having experienced the service, and in particular after relatively large shocks to their portfolios.

We start by computing, for each investor, the distance between the current allocation and the target allocation. For robo-takers, we define the target allocation as the one proposed by the robot and accepted by the investor. For robo-curious, we define the target allocation as the one held at the time of completion of the robo-survey, which the investor has preferred to the one proposed by the robot.

The robot is programmed to send email alerts if the distance between the current and the target allocation exceeds a given threshold. Several dimensions are considered, based on the proportion of assets allocated to different types of funds and on a synthetic measure of portfolio risk (SRRI). Accordingly, we construct a dummy equal to one if the distance is above the corresponding threshold in at least one dimension. That is,

$$Alert_{j,t} = \begin{cases} 1 & \text{if } max_z\{|\omega^z_{j,t} - \hat{\omega}^z_j| - \tau^z\} > 0, \\ 0 & \text{otherwise,} \end{cases} \tag{3}$$

where $\omega^z_{j,t}$ denotes the (value) proportion of funds of type z in contract j at time t, $\hat{\omega}^z_j$ is the proportion of funds of type z in contract j recommended by the robot and τ^z is a given threshold whose exact value is confidential. On average, in our sample, investors receive an alert in 8.6% of the months after the subscription.[14]

The variable $Alert_{j,t}$ can be constructed also for robo-curious, and it identifies the alerts that the robot would have sent had they taken the robot.[15] We can then measure, for robo-takers and robo-curious, how the distance between current and target equity exposure varies with the reception of the alert depending on whether or not the investor has accepted the robo-service.

We start by checking whether the reception of the alert is associated with an increased attention to the portfolio. In column 1 of Table 3, we observe that indeed upon reception of the alert investors are more likely to connect to the platform; the number of connections increases by 0.31 connections per month, relative to an increase of (statistically not-significant) 0.14 connections associated with the counterfactual alert.

We then analyze the associated rebalancing behaviors. In column 2, we consider

[14]The corresponding standard deviation is 28%, showing a significant variation in the number of alerts across investors and over time.

[15]Given our definition of target allocation, $Alert_{j,t}$ can only be constructed after the robo-adoption (for robo-takers) or its refusal (for robo-curious).

the probability of rebalancing upon reception of the alert (for robo-takers) or of the counterfactual alert (for robo-curious). The dependent variable is a dummy equal to one if the investor rebalances the portfolio in month t. We observe that robo-takers, who actually receive the alert, are 29% more likely to rebalance their portfolio, as compared to a baseline probability of rebalancing of 10% for robo-curious.[16]

As shown from the coefficient on *Alert*, robo-curious too are more likely to rebalance their portfolios in months in which they would have received an alert had they taken the robot. This is intuitive: alerts are typically engendered by sufficiently large market shocks, and these same shocks may induce robo-curious to rebalance their portfolio.[17] At the same time, the probability of rebalancing in those months is significantly larger for robo-takers, which highlights the distinctive role of the robot's alerts.

We then consider how by receiving the alert robo takers change their distance between the current and the target equity share. That is, if we denote by $x_{j,t}$ the equity share in contract j at time t and by τ_j the target equity share of contract j, our dependent variable is

$$|x_{j,t+1} - \tau_j| - |x_{j,t-1} - \tau_j|, \tag{4}$$

where $|.|$ is the absolute value and t is the month of the reception of the alert. We observe in column 3 that robo-takers decrease their distance by 4% more than robo-curious. The effect is large: conditionally on being alerted, the average distance is 6.2%.

An important observation is that, conditional on rebalancing, robo-curious investors rebalance in the opposite direction than robo-takers. While robo-takers tend to stay closer to their target equity exposure (i.e., they decrease equity exposure after good returns and increase it after bad returns), robo-curious investors remain further away from the target. This may reflect for example a tendency to be passive, or to chase the trends, increasing equity exposure after good returns and decreasing it after bad returns. Several studies document that trend-chasing behaviors may be harmful to portfolio returns (see e.g. Goetzmann and Kumar (2008), Greenwood and Nagel (2009), Bianchi (2018)) and, as we will show in the next section, these differences in rebalancing behaviors between robo-takers and robo-curious are an important dimension for understanding how the robot affects portfolio returns.

[16]Karlan, McConnell, Mullainathan and Zinman (2016) show that monthly reminders via SMS increase savings. Lee (2019) shows that overspending messages generated by a dedicated app induce users to reduce their spending.

[17]In this sense, robo-curious constitute a useful control group, as their portfolios can be viewed as similar to the ones that robo-takers would display were they not receiving the alerts.

In columns 4 and 5, we restrict to robo-takers and we compare the effect of our alert with another alert which investors receive if they have not completed the profiling survey as requested by the regulator (MIF). We observe that the effect of the MIF alert is in fact opposite (and much smaller) than the one of robo-alerts, confirming that the robot makes investors' portfolio closer to their target thanks to its specific alert.

An important observation is that the above results exploit behavioral changes occurring around the reception of alert, not around the robot's take-up. As such, it is unlikely that they are driven by unobservable shocks simultaneously affecting take-up and investment behaviors (discussed in more details in Section 5). In order to check that no confounding factors may occur around the reception of the alert, we verify that robo-takers and robo-curious follow parallel trends before the reception of the alert (for takers) or of the counterfactual alert (for curious). As shown in Figure 2, no significant pre-treatment differences can be observed in terms of level of attention, propensity to rebalance, and change in distance with the target allocation. This is reassuring on the view that the effects uncovered above are driven by the shocks induced by the robot's alerts.[18]

4 Risk-Taking and Returns

In this section, we first consider whether the changes in trading patterns described above are associated with changes in investors' exposure to risk. We then analyze the associated changes in portfolio returns, distinguishing between a static effect induced by a change in portfolio allocation occurring at the time of the robo-subscription and a dynamic effect induced by a change in the way in which investors rebalance their portfolios and modify their risk exposure over time.

Throughout this analysis, we use returns estimated directly from the liquidation value of the various funds, hence *net* of management and fund fees. Moreover, as mentioned, rebalancing activities are not subject to any transaction cost in our setting. We winsorize returns at the 0.5% level.

[18]From Figure 2 one may also notice the large confidence intervals at the time of the reception of the alerts, which is intuitive given the large heterogeneity in investors' responses to the alerts (see Section 5).

4.1 Equity Exposure and Portfolio Returns

We start by analyzing whether the robot's take up is associated with systematic changes in investors' equity exposure, which we compute as the value of equity over the total value of the portfolio, averaged over the previous 12 months. As shown in column 1 of Table 4, the robo-subscription is associated with an increase in the equity share by 3%, as compared to an average equity share of 22%.

We then investigate how this change in risk exposure translates into changes in portfolio returns. In column 2 of Table 4, we show that the robo-treatment is associated with an increase in returns by 2.6% per year. This effect is large, compared to an average return of 3%. In column 3, we control for the equity share, and observe that the robo-treatment is associated with an increase of 2.3% in yearly returns, which is again substantial.

These estimates should be interpreted with care, given that we are considering realized returns over a relatively short period of time. In order to further investigate their robustness, we consider how much of the effects on realized returns is driven by a change in exposure to standard risk factors. Following Reher and Sokolinski (2021), we consider a 5-factors model including 3 equity factors (Fama-French's market, size, value) and 2 fixed-income factors (Barclays' U.S. and Global Bond Index, taken from Bloomberg). We consider returns net of the U.S. risk-free rate, computed as the one-month Treasury yield (also taken from Ken French's library), and regress each fund's excess return over the U.S. risk-free rate to calculate the beta exposures of each fund. Regressions are based on the longest possible time-series, from 1990, when our fixed-income factor become available, to 2021, the end of our sample. As we are interested in highlighting the possibility dynamic effects on returns (in the next section), we consider time-varying expected returns. We define $R_t(x)$ as the return of each risky fund x (i.e., equity, balanced, bond, employer stock funds), in excess of the U.S. risk-free rate.[19] We compute the fund's return as the cross product of the fund's beta $\beta^f(x)$ and the realized returns of the corresponding factor R_t^f,

$$R_t(x) = \sum\nolimits_f \beta^f(x) R_t^f,$$

and the expected return of each portfolio based on each fund's portfolio weight.

We report our results in columns 4 and 5 of Table 4. We observe that subscribing to the robot is associated with an increase in expected returns by about 2.6% per year, as compared to an average expected return of 6%. Controlling for the equity share, the estimated increase in expected returns is equal to about 2.1% per year. These estimates

[19] For money market funds, we set these returns equal to the U.S. risk-free rate.

are essentially the same as those based on realized returns, which is reassuring on the robustness of our results. The extra returns generated by the robot cannot be attributed to "luck" but rather to the robot's *dynamic* exposure to standard risk factors.[20]

To have a rough measure of the euro value of these extra returns, an increase in returns of 2.6% would be associated with an increase in final wealth by about $19,043$ euros for an investor with average investment in the plan ($34,811$ euros) and average horizon (17 years).[21] While obviously rough, these figures suggest that the robot can have a significant impact on investors' wealth accumulation in the long run.

4.2 Static and Dynamic Effects

After having subscribed to the robot, investors may potentially change their portfolios in two dimensions. First, at the time of the subscription, they move from their current allocation to the one proposed by the robot, what we call a static effect. Second, investors may change the way in which they rebalance their portfolio over time, which we call a dynamic effect. In this section, we investigate how the two effects contribute to the observed changes in risk exposure and portfolio returns. Our objective is not to assess the optimality of the robot's rebalancing recommendations, which essentially aim at inducing constant portfolio weights, but rather to compare them to counterfactual rebalancing decisions investors would have taken without the robot.

In general, separating the static and the dynamic effects is challenging since we cannot directly observe the rebalancing behaviors (say, passive, contrarian, or trend chasing) the investor would have displayed without the robot; these rebalancing behaviors may vary considerably across investors and over time. This makes it hard to estimate the returns the investor would have experienced had she taken the robot at a given time t^* without changing her rebalancing behaviors at time $t > t^*$.

In our setting, however, we can exploit the knowledge of the robo-algorithm and identify a set of investors satisfying two conditions: i) they are robo-curious, i.e. they have completed the profiling survey proposed by the robot while eventually declining the service and ii) the portfolio allocation they hold when completing the survey is close to the one that the robot would have implemented. These investors, whom we call

[20] Indeed, assuming constant expected returns, we would estimate a lower increase in expected returns, ranging from 0.5 to 1% depending on the specifications. Focusing on time-varying expected returns allows us to better highlight the role of rebalancing, as we investigate next.

[21] In terms of fees associated with the service, this investor would pay around 257 euros in total over the 17 years.

curious close, are a useful reference point as their behavioral changes had they taken the robot would have been essentially those associated with rebalancing behaviors after subscription, since by definition the robot would not have implemented a large change to their initial allocation. Moreover, comparing the returns of curious close to those of robo-takers allows us to provide an estimate of the change in returns induced by a change in rebalancing behaviors associated with the robo-service, what we call a dynamic effect.[22]

We compute, for each robo-curious, the distance in equity share between the allocation held when completing the robo-survey and the one recommended by the robot. In our baseline analysis, we define curious close investors as those within a 5% distance in equity share, which includes 1,295 investors and that corresponds to 7% of the population of robo-curious. In robustness checks, reported in the Online Appendix (Tables A2 and A3), we consider a 10% threshold in the distance (corresponding to 15% of the robo-curious) and a 15% threshold (corresponding to 21% of the robo-curious).

We repeat the analysis of Table 4 while using the sample of curious close as control; results are reported in Table 5. We notice in column 1 that the average increase in equity share of robo-takers relative to curious close is about 1%, which is much smaller than the one reported in Table 4 when comparing to the entire population of robo-curious. At the same time, when investigating the effects of the robot's alerts as in Table 3, we observe that curious close investors display significantly different rebalancing behaviors than robo-takers (and, in fact, similar to those of the entire population of robo-curious), confirming the view that the observed changes in rebalancing behaviors are associated with the robot's recommendations.[23]

In terms of portfolio returns, we show in column 2 that when comparing to curious close investors, the robo-service is associated with an increase in returns of about 1.2%. Compared to the effect when using the entire population of robo-curious, which is equal to 2.6% (Table 5, column 1), this corresponds to about 46% of the total increase. By construction, this effect is essentially driven by changes in rebalancing behaviors after subscription, rather than by a change in portfolio allocation at the time of the subscription. A similar pattern emerges from the estimates in column 3, in which we control for the equity share, and in columns 4-5, in which we consider expected returns. When comparing to curious close, the effect of the robot on returns is between 42 and 46% of

[22]We are grateful to an anonymous Referee for suggesting this approach.

[23]We report the results on alerts in the Online Appendix (Table A1). We notice they are remarkably similar to our baseline analysis, confirming the view that the effects on rebalancing are likely driven by the robot's alerts, rather than by possibly unobservable heterogeneity between the robo-curious and the robo-takers.

the total effect, highlighting that a key determinant of the increase in returns we observe is the dynamic effect induced by the way in which investors rebalance their portfolios over time.

One may wonder whether the difference in returns between robo-takers and curious close investors tend to occur under particular market conditions. In Figure 3, we plot how the difference in cumulative returns between robo-takers and curious close investors evolves over time, together with cumulative market returns. The figure does not reveal a clear pattern between the two time-series, suggesting that the difference in returns induced by the robot is not driven by specific market episodes.

It may be useful to put these estimates in perspective with other estimates of rebalancing premia. Comparing portfolio rebalancing with constant weights to a buy-and-hold strategy, Maeso and Martellini (2020) find an annualized rebalancing premium of about 1% in the U.S. stock market, controlling for several risk factors. Similarly, for a diversified portfolio composed only of stocks and bonds, Ang, Brandt and Denison (2014) estimate a rebalancing premium of 0.14 in terms of average returns over realized volatility. As average volatility in our setting is around 14%, this would correspond to a premium of 1.96%. These estimates confirm the general message that changing rebalancing behaviors can be a key determinant of portfolio performance.[24]

5 Counterfactual: Automatic Rebalancing

As stressed above, we view the possibility for investors to retain control over their rebalancing decisions as an important feature of our setting, potentially reducing algo-aversion and promoting financial capability. At the same time, evidence in other domains shows that having humans-in-the-loop may be harmful for performance (see for example Ge et al. (2021) and Green and Chen (2019)).

In this section, we investigate the potential financial costs of letting investors decide whether or not to follow the robot. We restrict our analysis to robo-takers and we consider a counterfactual scenario assuming the robot were able to automatically rebalance the investor's portfolio. We define the counterfactual returns investors would have experienced had they rebalanced their portfolio immediately upon reception of the alert and exactly as suggested by the robot. We consider how a given rebalancing decision affects

[24]Evidence along those lines also appears in the mutual fund industry, where according to Berk and Van Binsbergen (2015) half of the value added can be attributed to improved diversification and half to market timing.

the returns in the next month, and for each investor we compute the difference between experienced and counterfactual returns, which we take as a measure of the cost (in terms of foregone returns) of retaining control.

We notice that, on average, the cost of retaining control is not large. In annual terms, counterfactual returns are 0.054% larger than actual returns. To get a better sense of the importance of these costs over time, we compute the cumulative returns of robo-takers and the counterfactual returns they would have experienced if the robot was able to automatically rebalance their portfolio. As comparison, we also compute the cumulative returns experienced by curious close investors (as defined in the previous section). The difference between automatic rebalancing and robo-takers provides an estimate of the effect, in terms of cumulative returns, of letting investors decide whether or not to follow the robot; the difference between robo-takers and curious close provides an estimate of the effect of the robot driven by changes in investors' rebalancing behaviors (what we have called a dynamic effect in the previous section).

We report these cumulative returns in Figure 4. We normalize the value of the portfolio of each group of investors at the beginning of the sample (September 2017) to 100. We compute the average returns experienced by a given group of investor and plot the corresponding cumulative returns over time. We notice that the cumulative returns of automatic rebalancing tend to exceed those of robo-takers. In turn, the cumulative returns of robo-takers tend to exceed those of curious close. Moreover, the magnitude of the difference between automatic and takers is significantly smaller than the one between takers and curious close. By the end of the sample (June 2021), the average cumulative returns induced by automatic rebalancing are about 12.85%, while for robo-takers they are about 12.61%, and for curious close about 7.03%.

While the average cost of retaining control is not large, we observe an important heterogeneity across investors. For an investor in the top 1% of the cost distribution, annual counterfactual returns are 4.1% larger than actual returns. The difference between counterfactual and actual returns is equal to 2.9% for an investor in the top 5%, to 0.7% for an investor in the top 10%, and to 0.4% for an investor in the top 25%. For an investor in the bottom 1% of the cost distribution, counterfactual returns are 8.7% smaller than actual returns. The difference between counterfactual and actual returns is equal to 1.3% for an investor in the bottom 5%, to 1% for an investor in the bottom 10%, and to 0.3% for an investor in the bottom 25%.

In order to further investigate the determinants of this important heterogeneity, we consider whether observable investors' characteristics and market conditions affect the propensity to rebalance upon reception of the alert and the probability that, conditional on rebalancing, the investor follows the robot's recommendation. In columns 1-3 of Table 6, we restrict our analysis to robo-takers in months of reception of the alert, and the dependent variable is a dummy equal to one if the investor rebalances her portfolio. We observe in column 1 that male, older and richer investors are more likely to rebalance on alert, as well as investors with lower past equity exposure. Overall, however, the magnitudes of these effects are small, compared to an average probability of rebalancing upon alert of about 44%.

In column 2, we explore the effect of market conditions, and consider in particular investors' behaviors during the bear markets between October and December 2018 and between February and March 2020. In these instances, investors were significantly less likely to rebalance on alert. For the average investor, the probability of rebalancing on alert during a bear market is 22.5%, as compared to 48% in other periods. These findings are confirmed in column 3, where we consider our regressors jointly.

We then analyze the probability that, conditional on rebalancing, the investor follows the robot's recommendation. In columns 4-6 of Table 6, we restrict to robo-takers in months at which they rebalance their portfolio, and the dependent variable is a dummy equal to one if rebalancing occurs as suggested by the robot. We observe that the average probability of following the robot is very large, about 85%. Interestingly, while as shown male investors are more likely to rebalance conditional on receiving an alert, they are also less likely to follow the robot conditional on rebalancing. These patterns are reminiscent of other studies in the literature showing that men tend to be less averse to algorithmic advice (see e.g. Niszczota and Kaszás (2020)), while at the same time being less likely to follow the advice, perhaps due to greater self-confidence (e.g. Barber and Odean (2001), Lusardi and Mitchell (2008)). In terms of other investors' characteristics, older and poorer investors are more likely to follow the robot, as well as investors with lower equity shares and larger past returns. The propensity to follow the robo does not change significantly in bear markets, and overall, these effects tend to be small relative to the large baseline probability.

Taken together, this evidence suggests that the costs in terms of foregone earnings of having investors-in-the-loop, as opposed to implementing an automatic rebalancing, are on average not large. At the same time, our figures reveal an important heterogeneity across investors. This heterogeneity is not significantly explained by investors' demographic characteristics; market conditions (and specifically investors' reluctance to rebalance during bear markets) appear more important.

6 Self-Selection

The decision to take-up the robo-service is voluntary and it can be driven by possibly unobservable characteristics that may also affect our outcome variables. In our previous analysis, we have addressed this issue by controlling for time-invariant individual-specific characteristics in a standard diff-in-diff specification. A possible concern is that individual-specific shocks may simultaneously drive the robo-subscription and a change in trading behaviors.

We then investigate whether our results are sensitive to the choice of the control group. In the baseline analysis, we have compared robo-takers to observationally similar individuals who have initiated the profiling process while eventually not subscribing to the robot. This comparison allows to condition on the (possibly time-varying) characteristics that may induce investors to express interest in the robo-service and compare the effects of the take-up relative to just observing the robot's profiling and recommendation. Moreover, while robo-curious could in principle replicate the robot's recommendation without subscribing, our results are clearly associated with the adoption of the robo-service, not just to the observation of the robo-recommendation. Alternatively, one may compare robo-takers to individuals who have been offered the service and did not express interest in the robot, which conditions on (possibly unobservable) factors correlated to having being exposed to the robot, or to individuals who have not been offered the service, noticing that the possibility to access the robot depends on a decision of the employer, not of the individual investor, and as such it may not be deriving from individual-level selection.

Considering alternative control groups is useful as unobserved differences may vary depending on whether we compare robo-takers to robo-curious or we use instead not-takers or not-exposed investors as comparison. While the nature of this heterogeneity is a priori not obvious, verifying the robustness of our findings when we vary the control group allows to make sure that our estimates are not driven by these unobserved differences.

In Table 7, we revisit our main results which exploit the robot's take-up as treatment. Our analysis of alerts (Table 3) focuses on how investors change their behaviors around the reception of alert; hence, as mentioned, it is less likely to be contaminated by individual-specific shocks occurring together with the robo-subscription. We consider the effects on attention (as in column 1 of Table 2), on trading activities (as in column 4 of Table 2), on equity exposure (as in column 1 of Table 4), and on returns (as in column 1 of Table 5). Our specification is the same as in Equation (1), except that the control group are those exposed to the robot but not-takers (columns 1, 3, 5, 7) or the not-exposed (columns 2, 4, 6, 8). Results are overall very similar to our baseline estimates, showing that the exact specification of the control group is not a key driver of our results. This also shows that our estimates are mainly driven by changes in behaviors within the group of robo-takers, rather than between the various groups of investors.

We perform additional robustness tests, reported in the Online Appendix (Table A4). We consider 2SLS regressions in which we instrument the robot's take-up of investor i at time t by the fraction of employees (excluding the investor at hand) working in the same firm as i that have adopted the robot at time t. Interactions on the workplace may be an important determinant of take-up, as this is partly driven by peer effects, or by some word of mouth learning about the service. The validity of the instrument does *not* require that firms' characteristics are orthogonal to take-up rates, nor that we abstract from firm-specific shocks that may also affect take-up rates. At the same time, however, the validity can be challenged as peer interactions can have direct effects on portfolio decisions, beyond those related to the robot's take-up (Maturana and Nickerson (2019), Ouimet and Tate (2020)). With this important caveat in mind, we notice that the instrument is a strong predictor of the propensity to take-up, and the estimated effects are again consistent with the baseline results.

7 Conclusion

We have found that having access to a robo-advisor induces investors to pay more attention to their portfolios, increase their trading activities and their exposure to risk. We have shown that an important dimension of these effects comes from the dynamic interaction with the robot. Following the robot's alerts, investors change their rebalancing behaviors so as to stay closer to their target allocation and they tend to increase their exposure to risk when subsequent risky returns tend to be larger, which results in larger portfolio returns.

Our analysis highlights the role of human-robot interactions (e.g., through the alerts) and more generally the importance of having investors being the ultimate decision makers on their portfolios, as opposed to fully delegating to the robot. Potentially, this aspect is key to improve investors' financial capabilities. In this way, rather than reducing investors' attention and awareness, the robo-service would become a tool to promote financial education, which we believe is a key aspect when assessing the long-run consequences of robo-advising. We view our analysis as a first step, we hope it can motivate further work in this promising direction.

References

Ang, A., Brandt, M. W. and Denison, D. F. (2014), 'Review of the active management of the norwegian government pension fund global', *Report to the Norwegian Ministry of Finance* .

Angelova, V., Dobbie, W. and Yang, C. S. (2022), 'Algorithmic recommendations and human discretion'.

Bansal, G., Nushi, B., Kamar, E., Horvitz, E. and Weld, D. S. (2021), Is the most accurate ai the best teammate? optimizing ai for teamwork, *in* 'Proceedings of the AAAI Conference on Artificial Intelligence', Vol. 35, pp. 11405–11414.

Barber, B. M. and Odean, T. (2001), 'Boys will be boys: Gender, overconfidence, and common stock investment', *Quarterly Journal of Economics* **116**(1), 261–292.

Berk, J. B. and Van Binsbergen, J. H. (2015), 'Measuring skill in the mutual fund industry', *Journal of Financial Economics* **118**(1), 1–20.

Bianchi, M. (2018), 'Financial literacy and portfolio dynamics', *Journal of Finance* **73**(2), 831–859.

Bianchi, M. and Brière, M. (2022), Robo-advising: Less ai and more xai?, *in* 'Machine Learning in Financial Markets: A Guide to Contemporary Practice', Cambridge University Press.

Brynjolfsson, E. (2022), 'The turing trap: The promise & peril of human-like artificial intelligence', *arXiv preprint arXiv:2201.04200* .

Burton, J. W., Stein, M.-K. and Jensen, T. B. (2020), 'A systematic review of algorithm aversion in augmented decision making', *Journal of Behavioral Decision Making* **33**(2), 220–239.

Calvet, L. E., Campbell, J. Y. and Sodini, P. (2009), 'Fight or flight? portfolio rebalancing by individual investors', *Quarterly Journal of Economics* **124**(1), 301–348.

Choukhmane, T. (2019), 'Default options and retirement saving dynamics'.

D'Acunto, F., Prabhala, N. and Rossi, A. G. (2019), 'The promises and pitfalls of robo-advising', *Review of Financial Studies* **32**(5), 1983–2020.

D'Acunto, F. and Rossi, A. G. (2020), 'Robo-advising', *CESifo Working Paper 8225* .

D'Acunto, F. and Rossi, A. G. (2023), 'Robo-advice: Transforming households into rational economic agents', *Forthcoming, Annual Review of Financial Economics* .

Dietvorst, B. J., Simmons, J. P. and Massey, C. (2018), 'Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them', *Management Science* **64**(3), 1155–1170.

Gargano, A. and Rossi, A. G. (2018), 'Does it pay to pay attention?', *Review of Financial Studies* **31**(12), 4595–4649.

Ge, R., Zheng, Z., Tian, X. and Liao, L. (2021), 'Human–robot interaction: when investors adjust the usage of robo-advisors in peer-to-peer lending', *Information Systems Research* **32**(3), 774–785.

Goetzmann, W. N. and Kumar, A. (2008), 'Equity portfolio diversification', *Review of Finance* **12**(3), 433–463.

Gomes, F., Haliassos, M. and Ramadorai, T. (2021), 'Household finance', *Journal of Economic Literature* **59**(3), 919–1000.

Green, B. and Chen, Y. (2019), Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments, *in* 'Proceedings of the conference on fairness, accountability, and transparency', pp. 90–99.

Greenwood, R. and Nagel, S. (2009), 'Inexperienced investors and bubbles', *Journal of Financial Economics* **93**(2), 239–258.

Karlan, D., McConnell, M., Mullainathan, S. and Zinman, J. (2016), 'Getting to the top of mind: How reminders increase saving', *Management Science* **62**(12), 3393–3411.

Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J. and Mullainathan, S. (2018), 'Human decisions and machine predictions', *Quarterly Journal of Economics* **133**(1), 237–293.

Lee, S. K. (2019), 'Fintech nudges: Overspending messages and personal finance management', *NYU Stern School of Business* .

Loos, B., Previtero, A., Scheurle, S. and Hackethal, A. (2020), Robo-advisers and investor behavior, Technical report, Working paper, University of Technology of Sidney.

Lusardi, A. and Mitchell, O. S. (2008), 'Planning and financial literacy: How do women fare?', *American economic review* **98**(2), 413–417.

Madrian, B. C. and Shea, D. F. (2001), 'The power of suggestion: Inertia in 401 (k) participation and savings behavior', *Quarterly Journal of Economics* **116**(4), 1149–1187.

Maeso, J.-M. and Martellini, L. (2020), 'Measuring portfolio rebalancing benefits in equity markets', *The Journal of Portfolio Management* **46**(4), 94–109.

Maturana, G. and Nickerson, J. (2019), 'Teachers teaching teachers: The role of workplace peer effects in financial decisions', *Review of Financial Studies* **32**(10), 3920–3957.

Mozannar, H. and Sontag, D. (2020), Consistent estimators for learning to defer to an expert, *in* 'International Conference on Machine Learning', PMLR, pp. 7076–7087.

Niszczota, P. and Kaszás, D. (2020), 'Robo-investment aversion', *Plos one* **15**(9), e0239277.

Noy, S. and Zhang, W. (2023), 'Experimental evidence on the productivity effects of generative artificial intelligence', *Available at SSRN 4375283* .

Ouimet, P. and Tate, G. (2020), 'Learning from coworkers: Peer effects on individual investment decisions', *Journal of Finance* **75**(1), 133–172.

Raghu, M., Blumer, K., Corrado, G., Kleinberg, J., Obermeyer, Z. and Mullainathan, S. (2019), 'The algorithmic automation problem: Prediction, triage, and human effort', *arXiv preprint arXiv:1903.12220* .

Reher, M. and Sokolinski, S. (2021), 'Automation and inequality in wealth management', *Available at SSRN 3515707* .

Reher, M. and Sun, C. (2019), 'Automated financial management: Diversification and account size flexibility', *Journal of Investment Management* **17**(2), 1–13.

Rossi, A. and Utkus, S. (2020), 'Who benefits from robo-advising? evidence from machine learning', *Available at SSRN 3552671* .

Seru, A., Shumway, T. and Stoffman, N. (2010), 'Learning by trading', *Review of Financial Studies* **23**(2), 705–739.

Siddarth, D., Acemoglu, D., Allen, D., Crawford, K., Evans, J., Jordan, M. and Weyl, E. G. (2021), 'How ai fails us', *arXiv preprint arXiv:2201.04200* .

Tables and Figures

Table 1: Descriptive Statistics

	p5	mean	p95	sd	N
Panel A: Statistics by Individual					
Takers					
Age	27.00	44.76	62.00	11.01	14,191
Female	0.00	0.28	1.00	0.45	14,191
Saving plan value	128.48	5,090	20,603	11,222	14,191
Total account value	183.45	14,145	63,520	31,584	14,191
Yearly variable remuneration	0.00	1,504	6,553	2,777	14,191
Number of saving vehicles	1.00	2.52	6.82	1.99	14,191
Number of connections per month	0.00	0.81	2.83	2.54	14,191
Equity share	0.00	0.16	0.54	0.19	14,141
Annual return (realized)	0.00	0.06	0.18	0.08	12,153
Curious					
Age	28.00	45.92	64.00	11.53	18,848
Female	0.00	0.36	1.00	0.48	18,848
Saving plan value	190.55	8,504	33,035	18,745	18,848
Total account value	293.53	21,010	87,841	41,752	18,848
Yearly variable remuneration	0.00	1,928	7,417	3,049	18,848
Number of saving vehicles	1.00	2.55	6.25	1.95	18,848
Number of connections per month	0.00	0.89	3.25	2.45	17,069
Equity share	0.00	0.16	0.54	0.19	18,793
Annual return (realized)	0.00	0.06	0.21	0.08	17,673
Not Takers					
Age	27.00	45.99	68.00	12.87	19,718
Female	0.00	0.42	1.00	0.49	19,707
Saving plan value	31.51	5,425	23,519	15,486	19,718
Total account value	46.59	11,263	50,938	33,603	19,718
Yearly variable remuneration	0.00	917	4,631	2,288	19,718
Number of saving vehicles	1.00	2.08	5.00	1.55	19,718
Number of connections per month	0.00	0.30	1.17	1.28	19,718
Equity share	0.00	0.12	0.50	0.19	19,640
Annual return (realized)	0.00	0.05	0.20	0.08	18,533
Not Exposed					
Age	28.00	48.20	70.00	13.12	19,188
Female	0.00	0.32	1.00	0.47	18,917
Saving plan value	60.73	8,592	36,036	27,505	19,189
Total account value	85.05	20,247	89,422	51,907	19,189
Yearly variable remuneration	0.00	900	4,502	2,035	19,189
Number of saving vehicles	1.00	2.29	6.00	2.03	19,189
Number of connections per month	0.00	0.42	1.67	1.66	19,189
Equity share	0.00	0.12	0.51	0.20	19,144
Annual return (realized)	-0.07	0.06	0.25	0.12	17,616

Continues on next page.

Table 1: Descriptive Statistics (continued)

	p5	mean	p95	sd	N
Panel B: Statistics by Contract and Month					
Age	30.00	48.36	65.00	11.18	4,795,438
Female	0.00	0.36	1.00	0.48	4,795,438
Saving plan value	0.00	9,130	41,752	25,209	4,795,438
Total account value	140	34,811	139,757	61,821	4,795,438
Yearly variable remuneration	0.00	2,405	9,270	3,690	4,795,438
Number of saving vehicles	1.00	4.21	10.00	2.87	4,795,438
Number of connections per month	0.00	1.01	4.00	3.17	1,697,422
Number of asset allocation changes	0.00	0.05	0.00	0.29	4,795,438
Number of asset allocation changes (robo)	0.00	0.01	0.00	0.11	4,795,438
Number of asset allocation changes (indiv)	0.00	0.02	0.00	0.23	4,795,438
Annual return (realized)	-0.09	0.03	0.24	0.13	3,174,911
Annual return (expected)	-0.01	0.06	0.25	0.10	3,173,832
Equity share	0.00	0.22	0.80	0.26	2,782,081

NOTE: This table reports descriptive statistics of our variables. In Panel A, we report statistics at the individual level for the various samples of investors. For each investor, we consider the average value of the variable before the introduction of the robo-service, between January 2016 and August 2017. In Panel B, we report all the observations by contract and over time for robo-takers and curious, which are used in our main analysis. Saving plan value refers to the single saving contract. Total account value is the aggregate across all contracts held by the same investor.

Table 2: Investors' Attention and Trading

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Variable	Number of Connections			Trades	Robo(>t)	Individual
Robo treated*after	0.283***	0.270***	0.140***	0.094***	0.044***	0.003
	(0.074)	(0.080)	(0.040)	(0.014)	(0.004)	(0.002)
Sample		No rem	No Sub			
Individual Fixed Effects	x	x	x			
Contract Fixed Effects				x	x	x
Time Fixed Effects	x	x	x	x	x	x
Observations	881,087	675,586	871,373	3,589,424	3,589,424	3,589,424
R-squared (within)	0.01	0.01	0.01	0.01	0.01	0.00

NOTE: This table reports the results of OLS regressions. In columns 1-3, the dependent variable is the number of connections per month. In column 2, we exclude the month before and the month at which the individual has received the variable remuneration. In column 3, the sample excludes the month of the robo-subscription. In column 4, the dependent variable is the number of allocation changes per month; in columns 5-6, the dependent variable is the number of allocation changes suggested by the robot and directly chosen by the individual, respectively. Controls include the average equity share and the average returns over the past 12 months, the account value in the previous month, the value of the yearly variable remuneration. Standard errors, double-clustered by individual and time (i.e., year-month), are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 3: Alerts and Rebalancing

	(1)	(2)	(3)	(4)	(5)
Dep. Variable	Connections	Rebalancer	Change in Distance Actual - Target Equity		
Robo treated*Alert	0.308***	0.295***	-0.046***		
	(0.051)	(0.048)	(0.004)		
Alert	0.144	0.104***	0.040***	-0.006***	
	(0.117)	(0.015)	(0.004)	(0.001)	
Alert MIF					0.001*
					(0.001)
Sample	Robo takers+curious			Robo takers	
Individual Fixed Effects	x				
Contract Fixed Effects		x	x	x	x
Time Fixed Effects	x	x	x	x	x
Observations	208,705	1,434,041	1,286,735	679,577	614,292
R-squared (within)	0.01	0.15	0.01	0.00	0.00

NOTE: This table reports the results of OLS regressions. In column 1, the dependent variable is the number of connections per month. In column 2, the dependent variable is a dummy equal to one if the investor rebalances the portfolio in month t. In columns 3-5, the dependent variable is the change in the distance between the actual and the target equity share between t+1 and t-1. In columns 1-3, the sample is restricted to robo-takers and robo-curious. Alert is a dummy equal to one if the distance between the actual and the target allocation is above the alert threshold, and to zero otherwise. For robo-takers, the target allocation is the one proposed by the robot; for robo-curious, it is the one held at the time of the completion of the robo-survey. In columns 4-5, the sample is restricted to robo-takers. Alert MIF is a dummy equal to one if the investor receives an alert as they have not completed the profiling survey requested by the regulator. Controls include the account value in the previous month, the value of the yearly variable remuneration, a dummy if the variable remuneration was received in the current month and a dummy if the variable remuneration was received in the past month. Standard errors, double-clustered by individual and time (i.e., year-month), are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4: Equity Exposure and Returns

	(1)	(2)	(3)	(4)	(5)
	Equity Share	Realized Returns		Expected Returns	
Robo treated*after	0.030***	0.026***	0.023***	0.026***	0.021***
	(0.003)	(0.005)	(0.005)	(0.003)	(0.003)
Equity Share			0.095***		0.136***
			(0.007)		(0.006)
Contract Fixed Effects	x	x	x	x	x
Time Fixed Effects	x	x	x	x	x
Observations	2,782,081	3,174,911	3,174,652	3,173,599	3,173,326
R-squared	0.01	0.00	0.01	0.01	0.05

NOTE: This table reports the results of OLS regressions. In column 1, the dependent variable is the equity share. In columns 2-4, the dependent variable is the annual returns at the contract level. In columns 3-5, the dependent variable is the expected annual returns at the contract level. Controls include the account value in the previous month, the value of the yearly variable remuneration, a dummy if the variable remuneration was received in the current month and a dummy if the variable remuneration was received in the past month. Standard errors, double-clustered by individual and time (i.e., year-month), are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 5: Equity Exposure and Returns: Comparing to Curious Close

	(1)	(2)	(3)	(4)	(5)
	Equity Share	Realized Returns		Expected Returns	
Robo treated*after	0.011**	0.012***	0.010***	0.011***	0.009***
	(0.004)	(0.003)	(0.003)	(0.002)	(0.002)
Equity Share			0.096***		0.139***
			(0.008)		(0.006)
Contract Fixed Effects	x	x	x	x	x
Time Fixed Effects	x	x	x	x	x
Observations	1,127,745	1,275,225	1,275,225	1,273,690	1,273,690
R-squared	0.00	0.01	0.02	0.01	0.08

NOTE: This table reports the results of OLS regressions. The control group is restricted to robo-curious for whom the difference between the equity share held at the time of the completion of the robo-survey and the one proposed by the robot was less than 5%. In column 1, the dependent variable is the equity share. In columns 2-4, the dependent variable is the annual returns at the contract level. In columns 3-5, the dependent variable is the expected annual returns at the contract level. Controls include the account value in the previous month, the value of the yearly variable remuneration, a dummy if the variable remuneration was received in the current month and a dummy if the variable remuneration was received in the past month. Standard errors, double-clustered by individual and time (i.e., year-month), are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 6: Robo Rebalancing

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Variable	Rebalancer on Alert			Robo Rebalancer		
Age	0.002***		0.002***	0.003***		0.003***
	(0.0002)		(0.0002)	(0.0001)		(0.0001)
Female	-0.015***		-0.019***	0.017***		0.021***
	(0.004)		(0.004)	(0.003)		(0.003)
Total account value (ln)	0.013***		0.009***	-0.019***		-0.034***
	(0.001)		(0.001)	(0.001)		(0.001)
Variable remuneration (10k)	-0.001		0.000	-0.007***		-0.004***
	(0.001)		(0.001)	(0.001)		(0.000)
Past equity share	-0.046***		-0.054***	-0.051***		-0.101***
	(0.009)		(0.009)	(0.006)		(0.006)
Past returns	0.003		-0.367***	0.153**		-0.147**
	(0.106)		(0.108)	(0.066)		(0.070)
Bear market		-0.257***	-0.257***		-0.001	0.007
		(0.005)	(0.005)		(0.005)	(0.005)
Time FE	x			x		
Observations	70,358	70,579	70,358	62,453	63,052	62,453
R-squared (within)	0.01	0.03	0.00	0.03	0.00	0.04

NOTE: This table reports the results of OLS regressions. In columns 1-3, the sample is restricted to robo-takers and to months in which they receive an alert, and the dependent variable is a dummy equal to one if the investor rebalances her portfolio. In columns 4-6, the sample is restricted to robo-takers and to months in which some rebalancing activity occurs, and the dependent variable is a dummy equal to one if the investor rebalances her portfolio following the robot's recommendation. Bear Market is a dummy equal to one between October and December 2018 and between February and March 2020. Standard errors, double-clustered by individual and time (i.e., year-month), are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 7: Control Groups

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. Variable	Connections		Trades		Equity		Returns	
Treated*after	0.270**	0.299***	0.093***	0.090***	0.046***	0.051***	0.029***	0.042***
	(0.097)	(0.090)	(0.018)	(0.016)	(0.004)	(0.005)	(0.005)	(0.008)
Control group	Exposed	Not-Exp	Exposed	Not-Exp	Exposed	Not-Exp	Exposed	Not-Exp
Individual FE	x	x						
Contract FE			x	x	x	x	x	x
Time FE	x	x	x	x	x	x	x	x
Observations	832,283	808,816	3,676,837	3,835,642	2,781,104	2,851,673	3,202,133	3,300,663
R-squared	0.01	0.01	0.01	0.01	0.02	0.02	0.01	0.01

Note: This table reports the results of OLS regressions. In columns 1 and 2, the dependent variable is the number of connections per month; in columns 3 and 4, the dependent variable is the number of allocation changes per month; in columns 5 and 6, the dependent variable is the equity share; in columns 7 and 8, the dependent variable is variable is the annual return. In columns 1,3,5 and 7, the control group are exposed individuals who did not take the robot. In columns 2,4,6 and 8, the control group are individuals who have not been offered the robo-service. Controls include the account value in the previous month, the value of the yearly variable remuneration, a dummy if the variable remuneration was received in the current month and a dummy if the variable remuneration was received in the past month. Standard errors, double-clustered by individual and time (i.e., year-month), are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Figure 1: Robo-Subscription, Investors' Behaviors and Returns

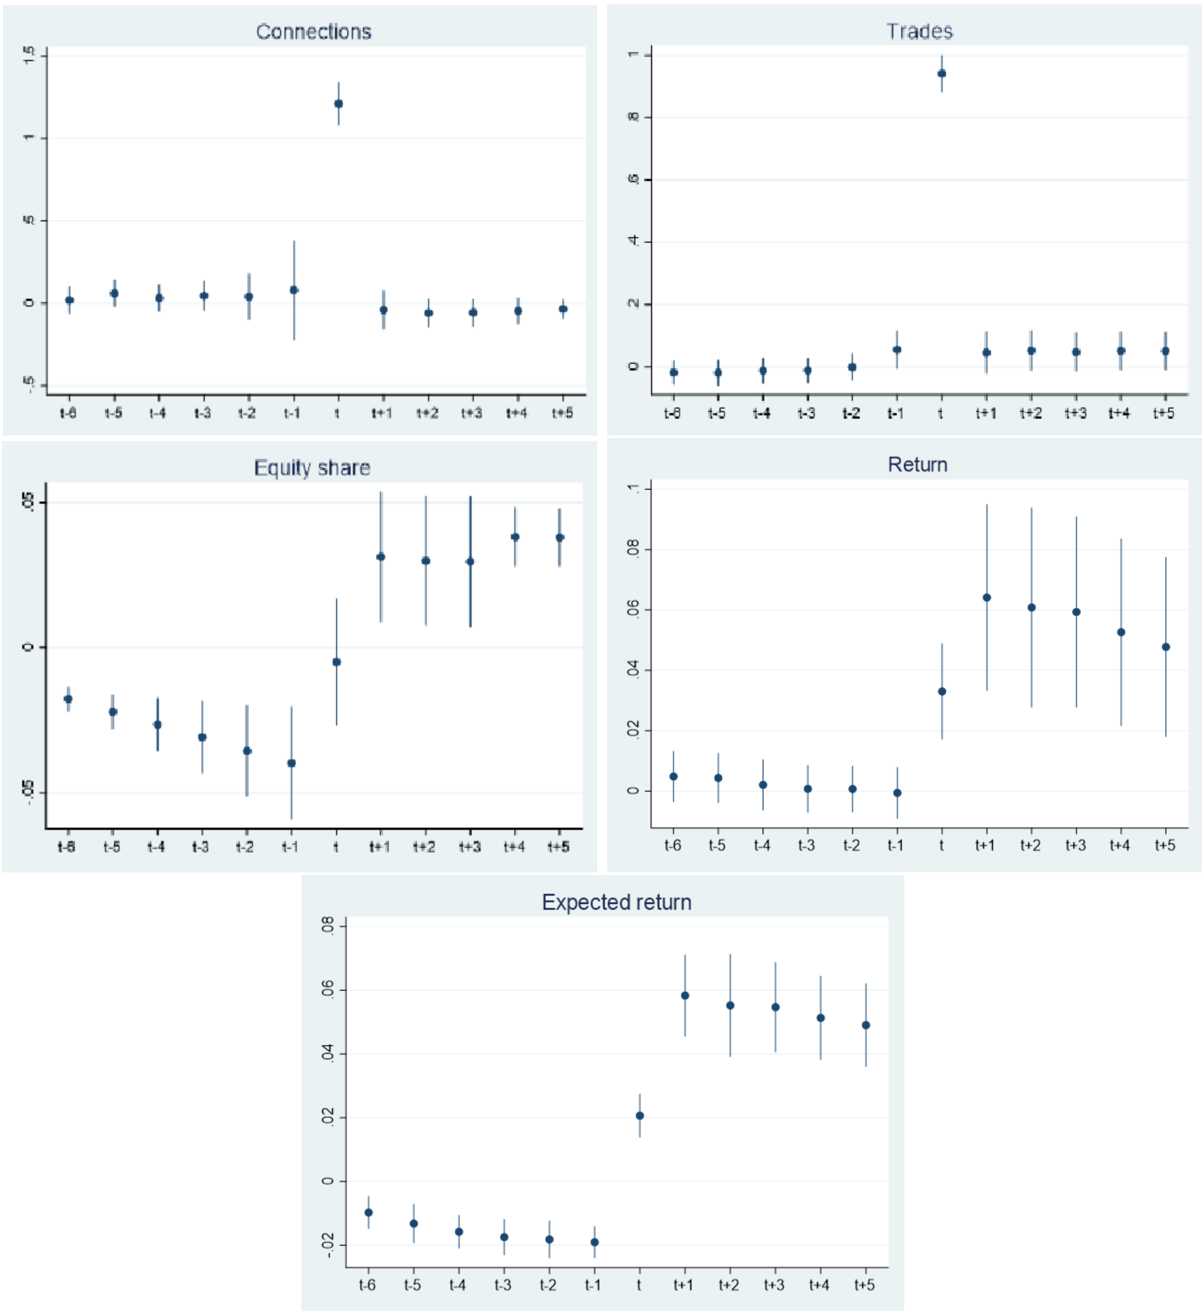

NOTE: This figure displays how the number of connections (top left panel), number of trades (top right panel), equity share (middle left panel), realized returns (middle right panel), and expected returns (bottom panel) vary around the subscription of the robot, as compared to the corresponding changes occurring for robo-curious. On the horizontal axes, t-6/t-1 correspond to months before the robo-subscription, t+1/t+5 correspond to months after the robo-subscription. The points correspond to the estimated coefficients, the bars correspond to 95% confidence intervals.

Figure 2: Alerts and Investors' Behaviors

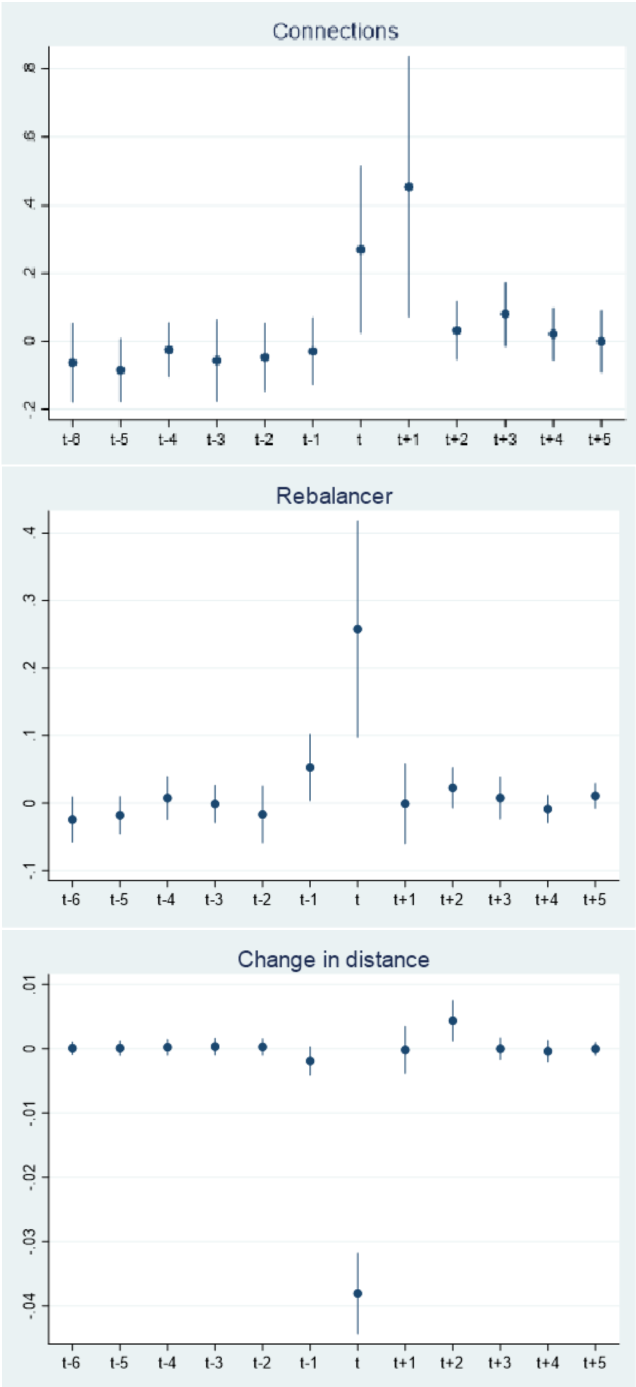

Note: This figure displays how the number of connections (top panel), the probability of rebalancing (middle panel), and the change in distance with the target allocation (bottom panel) vary around the reception of the alert for robo-takers, as compared to the corresponding changes occurring around the counterfactual alert for robo-curious. On the horizontal axes, t-6/t-1 correspond to months before the alert, t+1/t+5 correspond to months after the alert. The points correspond to the estimated coefficients, the bars correspond to 95% confidence intervals.

Figure 3: Cumulative Returns Differences and Market Returns

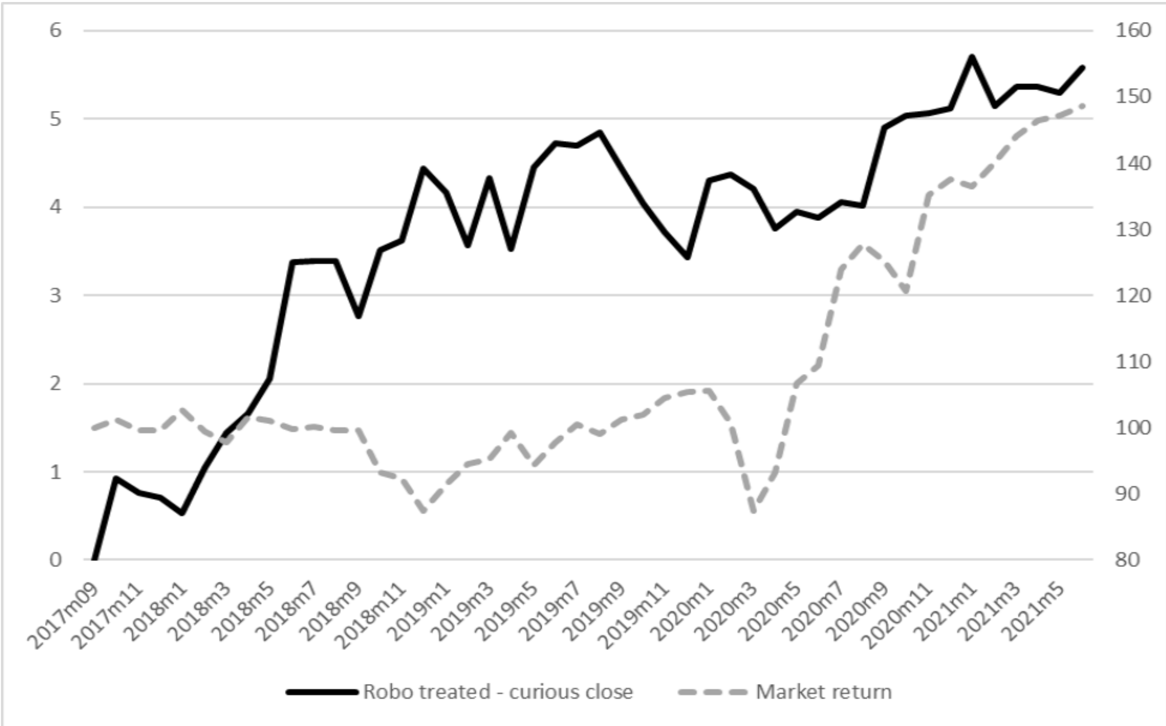

NOTE: In this figure, the solid line displays the difference in cumulative returns between robo-takers and curious-close investors over time. Curious-close investors correspond to robo-curious for whom the difference between the equity share held at the time of the completion of the robo-survey and the one proposed by the robot was less than 5%. Cumulative returns are computed as the average returns experienced by a given group of investor; the value of the portfolio at the introduction of the robot (September 2017) is normalized to 100 for all groups of investors. The scale of cumulative return differences is reported on the left part of the vertical axis, corresponding to return differences in percentage points. The grey dashed line displays how cumulative market returns evolve over time. Market returns at time t are defined as the unweighted average returns of the risky funds across all investors at time t. Market returns at the introduction of the robot (September 2017) is normalized to 100, and the scale is reported on the right part of the vertical axis. On the horizontal axis, time is expressed in months.

Figure 4: Cumulative Returns: Automatic Rebalancing, Robo-Takers and Robo-Curious

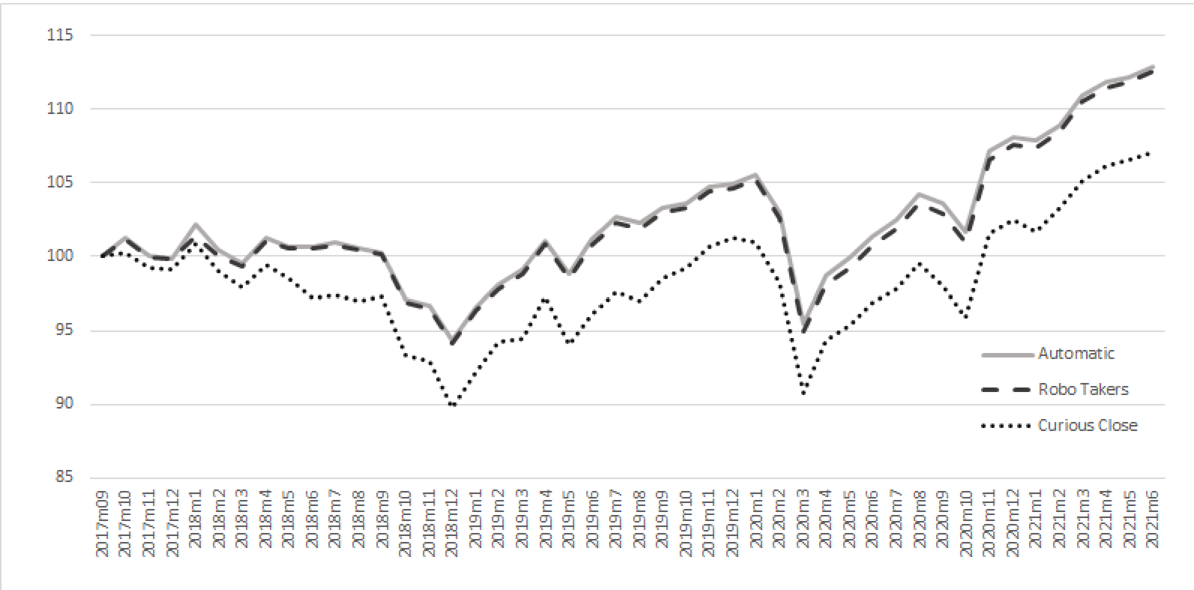

Note: This figure displays the cumulative returns experienced by various groups of investors over time. The value of the portfolio at the introduction of the robot (September 2017) is normalized to 100 for all groups of investors. On the horizontal axis, time is expressed in months; on the vertical axis, cumulative returns are computed as the average returns experienced by a given group of investor. The dark-grey dotted line correspond to robo-curious for whom the difference between the equity share held at the time of the completion of the robo-survey and the one proposed by the robot was less than 5%. The middle-grey dashed line corresponds to robo-takers. The light-grey solid line corresponds to (fictitious) investors who would automatically rebalance their portfolio immediately upon reception of the alert and exactly as suggested by the robot.

Online Appendix

Table A1: Rebalancing and Curious Close

	(1)	(2)	(3)
Dep. Variable	Connections	Rebalancer	Change in Distance
Robo treated*Alert	0.329***	0.309***	-0.035***
	(0.077)	(0.044)	(0.003)
Alert	0.130	0.096***	0.029***
	(0.101)	(0.018)	(0.004)
Individual Fixed Effects	x		
Contract Fixed Effects		x	x
Time Fixed Effects	x	x	x
Observations	124,835	904,441	817,907
R-squared (within)	0.01	0.17	0.01

NOTE: This table reports the results of OLS regressions. The control group is restricted to robo-curious for whom the difference between the equity share held at the time of the completion of the robo-survey and the one proposed by the robot was less than 5%. In column 1, the dependent variable is the number of connections per month. In column 2, the dependent variable is a dummy equal to one if the investor rebalances the portfolio in month t. In columns 3, the dependent variable is the change in the distance between the actual and the target equity share between t+1 and t-1. Alert is a dummy equal to one if the distance between the actual and the target allocation is above the alert threshold, and to zero otherwise. For robo-takers, the target allocation is the one proposed by the robot; for robo-curious, it is the one held at the time of the completion of the robo-survey. Controls include the account value in the previous month, the value of the yearly variable remuneration, a dummy if the variable remuneration was received in the current month and a dummy if the variable remuneration was received in the past month. Standard errors, double-clustered by individual and time (i.e., year-month), are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table A2: Curious Close : 10% threshold

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. Variable	Connections	Rebalancer	Ch. Distance	Realized Returns		Expected Returns	
Robo treated*Alert	0.305***	0.316***	-0.034***				
	(0.070)	(0.043)	(0.003)				
Alert	0.169*	0.090***	0.028***				
	(0.095)	(0.019)	(0.004)				
Robo treated*after				0.018***	0.016***	0.016***	0.013***
				(0.004)	(0.003)	(0.003)	(0.002)
Equity Share					0.097***		0.140***
					(0.008)		(0.006)
Individual FE	x						
Contract FE		x	x	x	x	x	x
Time FE	x	x	x	x	x	x	x
Observations	117,525	858,336	776,206	1,447,895	1,447,895	1,446,362	1,446,362
R-squared (within)	0.01	0.18	0.01	0.01	0.02	0.01	0.07

NOTE: This table reports the results of OLS regressions. The control group is restricted to robo-curious for whom the difference between the equity share held at the time of the completion of the robo-survey and the one proposed by the robot was less than 10%. In column 1, the dependent variable is the number of connections per month. In column 2, the dependent variable is a dummy equal to one if the investor rebalances the portfolio in month t. In columns 3, the dependent variable is the change in the distance between the actual and the target equity share between t+1 and t-1. Alert is a dummy equal to one if the distance between the actual and the target allocation is above the alert threshold, and to zero otherwise. For robo-takers, the target allocation is the one proposed by the robot; for robo-curious, it is the one held at the time of the completion of the robo-survey. in columns 4-5, the dependent variable is the annual returns; in columns 6-7, the dependent variable is the expected annual returns. Controls include the account value in the previous month, the value of the yearly variable remuneration, a dummy if the variable remuneration was received in the current month and a dummy if the variable remuneration was received in the past month. Standard errors, double-clustered by individual and time (i.e., year-month), are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table A3: Curious Close : 15% threshold

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. Variable	Connections	Rebalancer	Ch. Distance	Realized Returns		Expected Returns	
Robo treated*Alert	0.329***	0.309***	-0.035***				
	(0.077)	(0.044)	(0.003)				
Alert	0.130	0.096***	0.029***				
	(0.101)	(0.018)	(0.004)				
Robo treated*after				0.021***	0.018***	0.019***	0.015***
				(0.004)	(0.004)	(0.003)	(0.003)
Equity Share					0.097***		0.140***
					(0.008)		(0.006)
Individual FE	x						
Contract FE		x	x	x	x	x	x
Time FE	x	x	x	x	x	x	x
Observations	124,835	904,441	817,907	1,583,096	1,583,096	1,581,592	1,581,592
R-squared (within)	0.01	0.17	0.01	0.01	0.01	0.01	0.07

NOTE: This table reports the results of OLS regressions. The control group is restricted to robo-curious for whom the difference between the equity share held at the time of the completion of the robo-survey and the one proposed by the robot was less than 15%. In column 1, the dependent variable is the number of connections per month. In column 2, the dependent variable is a dummy equal to one if the investor rebalances the portfolio in month t. In columns 3, the dependent variable is the change in the distance between the actual and the target equity share between t+1 and t-1. Alert is a dummy equal to one if the distance between the actual and the target allocation is above the alert threshold, and to zero otherwise. For robo-takers, the target allocation is the one proposed by the robot; for robo-curious, it is the one held at the time of the completion of the robo-survey. in columns 4-5, the dependent variable is the annual returns; in columns 6-7, the dependent variable is the expected annual returns. Controls include the account value in the previous month, the value of the yearly variable remuneration, a dummy if the variable remuneration was received in the current month and a dummy if the variable remuneration was received in the past month. Standard errors, double-clustered by individual and time (i.e., year-month), are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table A4: IV Estimates

	(1)	(2)	(3)	(4)
Dep. Variable	Connections	Trades	Equity	Returns
Robo treated*after	0.163**	0.067***	0.098***	0.062***
	(0.075)	(0.027)	(0.009)	(0.015)
	First Stage: Robo Treated			
Fraction of treated employees	14.309***	3.225***	2.519***	3.064***
	(1.918)	(0.330)	(0.345)	(0.362)
F-Stat (first stage)	55.67	95.55	53.21	71.66
Individual Fixed Effects	x			
Contract Fixed Effects		x	x	x
Time Fixed Effects	x	x	x	x
Observations	807,785	3,823,101	2,843,430	3,290,246
R-squared (centered)	0.011	0.006	0.005	0.008

Note: This table reports the results of 2SLS regressions in which the probability to adopt the robo-service is instrumented by the fraction of employees in the same firm who have taken-up the robot. In column 1, the dependent variable is the number of connections per month; in column 2, the dependent variable is the number of allocation changes per month; in column 3, the dependent variable is the share; in column 4, the dependent variable is the annual return. Controls include the account value in the previous month, the value of the yearly variable remuneration, a dummy if the variable remuneration was received in the current month and a dummy if the variable remuneration was received in the past month. Standard errors, clustered by firm and time (i.e., year-month), are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.