Human-Robot Interactions in Investment Decisions*

Milo Bianchi† Marie Brière‡

November 7, 2022

Abstract

We study the introduction of robo-advising on a large representative sample 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 to an increase in investors' attention and trading activities. In particular, following the robot's alerts, investors change their rebalancing behaviors so as to stay closer to their target allocation, which translates into higher risk-adjusted returns. We discuss how human-robot interactions can improve financial capability.

Keywords: Robo-Advising, Human-robot Interaction, Portfolio Dynamics.
JEL codes: G11; G51; G41; G23; D14.

*We are grateful to Francesco D'Acunto, Francesco D'Amuri, Laurent Bach, Pascal Blanqué, 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 comments. We also thank Xavier Collot, Corinne Laboureix, Olivier Melennec, Fabien Thevenot for insightful explanations on the institutional setting and Arnaud Delavoet, Marc-Antoine Hubert, Anthony Lapeyre and Mohamed Youssef for help in the data gathering process. 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 (Bianchi and Brière (2021)), and they adopt verifiable procedures, which may limit the extent of biased advice (Philippon (2019)). 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. 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.[1]

We investigate human-robot interactions by exploiting the introduction of a robo-advising service by a large French asset manager. The robot starts by eliciting information on the client's characteristics, builds the client's profile, and proposes a portfolio allocation. If the client accepts the proposal, the robot implements the allocation. Over time, the robot sends email alerts if the current portfolio allocation ends up being too far from the target allocation. These alerts suggest to connect to the platform and to rebalance the portfolio towards the target. The distinctive feature of this service is that the robot gives advice to the investors, both at the time of the subscription and over time, while leaving them free to follow or to ignore the advice. This makes it different from the more common robo-advisors (discussed in the literature review) 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.

A recurrent theme of our analysis is that these interactions are key to understanding the ultimate effects of the robot on financial outcomes. 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.[2] Second, these interactions allow us to study how the reliance on the robo-service evolves over time, say as investors experience market shocks or as new investment

[1] 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.

[2] In a forecasting task, Dietvorst, Simmons and Massey (2018) show that participants are more willing to rely on an automated advice when they can even slightly modify the algorithm. Similarly, Burton, Stein and Jensen (2020) present several experimental settings in which algorithm aversion is reduced when giving participants some control over the underlying algorithm. Bianchi and Brière (2021) review some evidence on finance applications.

opportunities arise. In these instances, 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. As we will show, our effects are largely driven by the way in which investors change their behaviors over time, suggesting that the robot can be used to improve investment decisions, while letting investors being the ultimate decision maker.

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. Absent the robot, employees self-manage their portfolios without any dedicated advice. We have access to account level data covering from September 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 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. We address this issue, first, by employing a diff-in-diff specification in which we compare changes in our outcome of interest associated to the robot's take-up to changes in a control group, which in our baseline analysis is defined by individuals who have not been exposed to the robo-service. We then consider alternative control groups (i.e., non-takers or curious) so as to isolate the effects of the robot from potentially confounding factors. Second, since the decision to take-up the robot may be influenced by interactions occurring at the workplace, we use the fraction of employees adopting the robot in a given firm as an instrument for the individual propensity to take-up. We provide more details of these alternative specifications, and show the robustness of our findings, as we proceed with the analysis.

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 0.8. Importantly, this effect persists even beyond the time of the robo-subscription.

This increase in attention is associated to 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. We show that the alerts sent by the robot when allocations get too far from the target are effective in inducing investors to rebalance their portfolio and stay closer to the target. Exploiting the knowledge of the algorithm governing the alerts, we 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 to rebalance, and it reduces the distance between current and target equity exposure by 4%, relative to an average distance of 11%.

These changes in trading behaviors have significant consequences for portfolio returns. We show that after subscribing to the robot, investors experience an increase in returns (net of fees) between 2% and 4% per year, controlling for various measures of portfolio risk. We decompose this increase in returns between a static effect occurring at the time of the subscription (say, as investors are moving closer to the efficient frontier) and a dynamic effect associated to the way in which investors rebalance their portfolios over time. We show that the static effect accounts for about 30% of the increase in returns, the most important part comes from a change in rebalancing behaviors.

Finally, we investigate the potential financial costs of letting investors decide whether or not to follow the robot, as opposed to completely delegating their portfolio rebalancing to the robot. Comparing the difference between the returns experienced by our investors and the counterfactual returns they would have experienced with automatic rebalancing, we observe a small average loss between 0.045% and 0.089% (depending on the horizon) in annualized returns and at the same time an important heterogeneity across investors depending on their rebalancing decisions.

Overall, these findings are encouraging on the possibility to promote human–robot interactions in the field of personal finance. We view our results as contributing 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 understand how the robot impacts investors' rebalancing behaviors and performance.[3]

The focus on human-robots interactions also distinguishes our paper from most other contributions, such as Reher and Sun (2019), Loos, Previtero, Scheurle and Hackethal (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 (2021) for overviews). Our paper shows that investors may improve their decisions even while retaining 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.[4]

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 client'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 client's profile (say, prudent, dynamic,..)

[3]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)).

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

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.[5]

The client 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 client accepts the proposal, the robot implements the allocation.[6] 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.[7] 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 14,576 individuals have subscribed to the robot and we refer to them as robo-takers. This corresponds to 17,069 accounts managed by the robot in 762 different firms. We observe no individual who subscribes to the robot and then terminates the service within our sample period.

We have extracted the trading records of all individuals who have taken up the robo-

[5]The robot is programmed to propose an allocation on the part of the portfolio which is not invested in employer's stock, which may have some specificities (e.g. in terms of matching rule) relative other stocks.

[6]Even if the client 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.

[7]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.

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 account in our sample period. We have followed these individuals up to June 2021, which gives us a panel covering the period September 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 37,764 euros, the median is 11,797 euros. These figures are comparable to those one can find in representative surveys.[8] 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).

3 Attention and Trading

In our baseline analysis, we explore the behavioral changes associated to 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 specifications at the contract level. We estimate the following equation:

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

where α_j and μ_t are contract and time fixed effects, $T_{j,t}$ 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.

[8] 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).

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 a sample of individuals who have not been exposed to the robot, while we consider alternative specifications in the robustness section. Summary statistics of the main variables used in the analysis are reported in Table 1.

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 and add individual, rather than contract, fixed effects. Moreover, our attention variables are only available on a shorter sample, from September 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 0.3 connections per month (the average is 0.8). Similar patterns hold with other measures of attention.[9]

One may question whether the increased attention is associated to 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 subscrip-

[9]We observe an increase in the number of minutes spent on the platform by 4.6 per month (the average is 6.2) and an increase in the number of web pages visited per month by 5.8 (the average is 6.5).

tion to the service. In column 3, we exclude the two months around 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.

Overall, these result show that investors do not take the robot as a substitute for their own attention. Rather, the robot is associated to an increased level of attention, which persists even beyond the time of its subscription.

3.2 Trading Activities

We now consider whether the increased level of attention is associated to 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 to 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 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, as defined at the time of the robo-subscription (or of the latest robo-profiling). In case of alert, the investor receives an email stating that there is discrepancy between the current and the target allocation, due to the investor's own trading or to a market shock, and she is suggested to connect to the dedicated website to consult her portfolio. The email alert is sent in the month at which the deviation occurs; if the deviation persists an additional email is sent the month after and then alerts stop, even if the deviation persists. 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 consider the sample of robo-takers and robo-curious (i.e., those individuals who have completed the robo-survey but have not subscribed to the service). For these investors, we can build 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).[10] Accordingly, we construct a dummy *Alert* that is equal to one if the distance is above the corresponding threshold in at least one dimension, and to zero otherwise. On average, in our sample, investors receive an alert in 8.6% of the months after the subscription.[11]

The variable *Alert* can be constructed also for robo-curious, and it identifies the alerts that the robot would have sent had they taken the robot.[12] 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

[10]The exact values of the thresholds are confidential.

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

[12]Given our definition of target allocation, *Alert* can only be constructed after the robo-adoption (for robo-takers) or its refusal (for robo-curious).

has accepted the robo-service.

We start by checking whether the reception of the alert is associated to 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 0.23 connections per month, relative to an increase of 0.11 connections associated to the counterfactual alert.[13] We then analyze the associated rebalancing behaviors. In column 2, we consider 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 or t+1, where t is the first month at which the distance between the actual and the target allocation exceeds the alert threshold. We observe that robo-takers, who actually receive the alert, are 28% more likely to rebalance their portfolio, as compared to a baseline probability of rebalancing of 13% for robo-curious.[14] In column 3, the dependent variable is the change in the distance between the actual and the target equity share, and we observe that robo-takers decrease their distance by 4% more than robo-curious. The effect is large: conditionally on being alerted, the average distance is 11%.

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.

4 Returns

We first consider whether the changes in trading patterns described above are associated to changes in portfolio returns, controlling for various measures of risk. We then decompose the total effect on returns between a static effect occurring at the time of the subscription and a dynamic effect associated to the way in which investors rebalance their portfolios over time. Finally, we compare investors' returns to counterfactual returns they would have obtained had they fully delegated portfolio rebalancing to the robot.

[13]Counterfactual alerts, just like actual alerts, occur after large changes in portfolio weights due to market shocks or active rebalancing, hence it is intuitive that even those alerts are associated to an increased attention.

[14]Karlan, McConnell, Mullainathan and Zinman (2016) show that monthly reminders via SMS increase savings.

4.1 Realized and Expected Returns

We start with the same specification as in (1), using realized returns as dependent variable. Throughout this analysis, we use returns net of management and fund fees, which we estimate directly from the liquidation value of the various funds. We winsorize returns at the 0.5% level. Results are presented in Table 4.

In column 1, we show that the robo-treatment is associated to an increase in returns by 4.2% per year. This effect is large, compared to an average return of 4%. We then consider various measures of portfolio risk. In column 2, we control for the equity share in the previous period; in column 3, we control for volatility, computed over a rolling window of 12 months; in column 4, we control for the beta of the portfolio, computed by taking as benchmark the returns of all the portfolios in our sample. We observe in these specifications that the robo-treatment is associated to an increase between 3.7 and 4.9% in yearly returns, which is again very large.

These estimates are crude and should be interpreted with care, also 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. We compute the beta of each fund based on the longest possible time-series, from 2020, 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, which we compute as the cross product of the fund's beta $\beta^f(x)$ and the realized returns of the corresponding factor R_t^f,

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

For money market funds, we set these returns equal to the U.S. risk-free rate. We can then compute the expected return of each portfolio based on each fund's portfolio weight.

We report our results in columns 5 and 6 of Table 4. We observe that subscribing to the robot is associated to an increase in expected returns by about 2.8% per year. Controlling for the past equity share, the estimated increase in expected returns is equal to about 2.1% per year.

To have a rough measure of the euro value of these extra returns, consider an investor with average investment in the plan (37,764 euros) and average horizon (17 years). An increase in yearly returns by 4.2% would be associated to an increase in final wealth by about 38,363 euros. Considering instead an increase in expected yearly returns by 2.8%, the associated increase in final wealth would be 23,127 euros. These extra returns can be compared to the fees associated to the subscription of the robot. On average, in our sample, investors pay a management fee equal to 0.01% of the amount invested in the saving plan. For robo-takers, the fee is on average equal to 0.05% of the portfolio.

Overall, these results suggest that the robot can have a significant impact on investors' wealth accumulation in the long run. In the next analysis, we investigate the determinants of the increase in returns associated to the robot by distinguishing a static effect occurring at the time of the subscription of the robot from a dynamic effect associated to different portfolio dynamics after the subscription.

4.2 Static and Dynamic Effects

As shown above, after subscribing to the robot, investors' portfolios change in two dimensions. First, at the time of the subscription, they move from their current allocation to the one proposed by the robot. We call this a static effect, which can positively impact returns to the extent that investors hold sub-optimal portfolio allocations (e.g. they wrongly estimate expected returns and risk or they choose allocations outside the efficient frontier). Second, investors may change the way in which they rebalance their portfolio over time, which we call a dynamic effect. The resulting impact on returns can be positive if for example investors tend to wrongly time the market. We investigate how the two effects contribute to the observed changes in portfolio returns.

Consider an investor who takes up the robot at time t^* and let us define $\omega_1(s,t)$ as the observed portfolio weight on asset s at time $t \geq t^*$ and $\omega_0(s,t)$ as the counterfactual weight on asset s the investor would have had without the robot. The associated portfolio returns at time t are $R_1(t) = \sum_s \omega_1(s,t-1)R(s,t)$, where $R(s,t)$ are the returns of asset s at time t, and the counterfactual returns without the robot are $R_0(t) = \sum_s \omega_0(s,t-1)R(s,t)$. According to the above estimates, the total effect $R_1(t) - R_0(t)$ is around 4.2% in yearly returns, and we wish to decompose this effect into a static and a dynamic effect. In general, this exercise in challenging since we cannot directly observe the returns the investor would have experienced had she taken the robot at time t^* without changing her rebalancing behaviors at time $t > t^*$. Moreover, these rebalancing behaviors (say, passive, contrarian, or trend chasing) may vary considerably across clients and over time.

In our setting, however, we can exploit the knowledge of the robo-algorithm. In our sample period, the robot's recommendations are essentially intended to induce constant portfolio weights.[15] Suppose that the robot were to keep the investor's current allocation unchanged and just change rebalancing behaviors and induce constant weights. The investor would then experience returns $C_0(t+1) = \sum_s \omega_0(s,t^*)R(s,t+1)$, where $\omega_0(s,t^*)$ are the portfolio weights observed just before the robot take-up at t^*,

$$\omega_0(s,t^*) = \frac{\omega_0(s,t^*-1)R(s,t)}{\sum_z \omega_0(z,t^*-1)R(z,t)}.$$

In this case, the robot would only have a dynamic effect, which can be computed as

$$D(t) = C_0(t) - R_0(t). \tag{2}$$

The static effect, due to the fact that the robot is also changing the investor's allocation at t^*, can be then computed as the residual

$$S(t) = R_1(t) - R_0(t) - (C_0(t) - R_0(t)). \tag{3}$$

We report our corresponding estimates in Table 5. In column 1, we estimate the static effect according to Equation (3) by considering the same diff-in-diff specification as in Equation (1) with $R_1(t) - C_0(t)$ as dependent variable. For robo-takers, we use the portfolio weights observed at the time of the robo-subscription; for investors who have not been exposed to the robot, we use the portfolio weights observed at the time of the first reception of the variable remuneration.

We observe that the static effect accounts for 1.3% of the total increase in returns, the remaining 2.9% is driven by the dynamic effect (the total effect, estimated in column 1 of Table 4, is 4.2%). In columns 2-4, we repeat the same decomposition controlling for various measures of risk, and find similar estimates in relative terms.

In columns 5-6 of Table 5, we repeat the same analysis considering instead expected returns (as in columns 5-6 of Table 4). We observe that the static effect accounts for about

[15] This would not be true over a longer time period, on which the robot would change the suggested allocations according to the investor's life-cycle.

1.1% of the increase in expected returns, the remaining 1.7% is driven by the dynamic effect (the total effect, estimated in column 5 of Table 4, is 2.8%). Similar results appear when we control for the equity share.

Overall, these figures show that a key determinant of the increase in returns we observe is given by a dynamic effect associated to the way in which investors rebalance their portfolios over time. According to our estimates, a change in rebalancing behaviors is associated to increase of about 290bps per year in terms of realized returns and of about 170bps in terms of expected returns. 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 100bp 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 196bp. These estimates confirm the general message that changing rebalancing behaviors can be a key determinant of portfolio performance.[16]

4.3 Automatic Rebalancing: Counterfactual Returns

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, Zheng, Tian and Liao (2021) on peer-to-peer lending and Green and Chen (2019) on judges' decisions).

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 construct a counterfactual scenario assuming the robot were able to automatically rebalance the investor's portfolio, and compute the associated returns. Our main variable of interest is the difference between the returns experienced by our investors and the counterfactual returns they would have experienced had they rebalanced their portfolio immediately upon reception of the alert and exactly as suggested by the robot.

We report our results in Table 6. In columns 1-3, we consider how a given rebalancing

[16]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.

decision affects the returns in the next month. We first show that, on average, the cost in terms of foregone returns of retaining control is not large. On average, in annual terms, counterfactual returns are 0.045% larger than actual returns (column 1).

We then explore whether we can detect a significant heterogeneity in this cost across investors, depending on their demographic characteristics and rebalancing behaviors. In column 2, we observe that demographic characteristics have little explanatory power on the difference between realized and counterfactual returns. In column 3, we consider the effects of different rebalancing behaviors, while adding fixed effects at the contract level (that also absorb time-invariant individual characteristics). We distinguish between those rebalancing activities that are freely initiated by the investor and those that follow the robot's recommendation, either immediately upon reception of the alert or at a later time. We observe that following the robot's recommendation is associated to an increase in returns, relative to counterfactual ones. The effect of free rebalancing is instead negative, its magnitude is large but not precisely estimated.

Since a given rebalancing decision may potentially impact returns for several months, we check whether the above effects are confirmed as we consider average returns (both actual and counterfactual) at different horizons. In columns 4-5, we consider a 3-month horizon; in columns 5-6, we consider a 6-month horizon.

We first notice that the average return loss becomes larger as we increase the horizon. Realized returns are 0.076% smaller than counterfactual returns if we consider average returns in the next 3 months, and the difference is 0.089% at a 6-month horizon. In terms of rebalancing behaviors, the same patterns uncovered above hold.

Overall, this analysis confirms the previous findings that the effects on rebalancing behaviors are a key driver for the overall effects of the robo-service in our setting. They also suggest 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.

5 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. In this section, we report a series of tests which aim at addressing this concern.

To simplify the exposition, all tables in this section have the same structure, which replicates our main results based on diff-in-diff specifications. We first consider the effect on attention (as in column 1 of Table 2) and on trading activities (as in column 4 of Table 2); we then consider returns (as in column 1 of Table 4) and the static change in returns (as in column 1 of Table 5).

5.1 Varying the Control Group

Our first test investigates the robustness of our findings when we vary the control group. In the baseline analysis, we have compared robo-takers to observationally similar individuals who have not been offered the service, so as to minimize biases deriving from individual-level selection. We now investigate the robustness of our findings when comparing robo-takers to individuals who have been offered the service and did not express interest in the robot or to robo-curious. In the first case, we condition on the exposure to the robot, and isolate the effect of taking up the service as opposed to not expressing interest. In the second case, we condition on the fact that the individual has expressed some interest in the robot, and compare the effect of the take-up relative to observing the robot's profiling and recommendation without subscribing to the service.

We report our results in Table 6, in which the control group are those exposed to the robot (columns 1,3,5,7) or the robo-curious (columns 2,4,6,8). In both cases, results are very similar to our baseline estimates. This is important since it shows that the exact specification of the control group is not a key driver of our results, our estimates are mainly driven by changes in behaviors within the group of robo-takers (as opposed to between groups). Moreover, while robo-curious could in principle replicate the robot's recommendation without subscribing, these results suggest that our estimated effect are associated to the adoption of the robo-service, not just to the observation of the robo-recommendation.

5.2 Instrumenting Take-up

As additional robustness check, we look for shifters to the propensity to take-up the service which are unlikely to be driven by individual-specific shocks. Interactions on the workplace may be an important determinant of take-up, which can be partly driven by peer effects, or by some word of mouth learning about the service. In fact, we observe an important variation in the take-up decisions across firms. In the 762 firms with at least one taker, the average take-up is 2.5% and the standard deviation is 8.1% with take-up ranging from 0.1% in the 5th percentile to 6.1% in the 95th percentile.

As intuitive, firms with low take-up may be different from firms with high take-up.[17] At the same time, 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. Rather, we require that these firm-level shocks are uncorrelated to shocks which are specific to the *individual* employee. Accordingly, we instrument the individual robot take-up at time t by the fraction of employees in the same firm that have adopted the robot at time t.

We report our results in Table 7. We observe that indeed the instrument is a strong predictor of the propensity to take-up. The estimated effects are once again very much in line with the baseline results.

6 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 it results in higher risk-adjusted returns. We have shown that an important dimension of these effects comes from the dynamic interaction with the robot, which is able to induce investors to rebalance their portfolio in a way that keeps them closer to the target allocation.

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 promote investors' learning on how to manage their portfolios and to improve their financial capabilities.[18] 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.

[17]Specifically, the fraction of treated individuals is positively associated to employees' average age, assets in the plans, and variable remuneration.

[18]Seru, Shumway and Stoffman (2010) study the dynamics of learning by trading; Loos et al. (2020) provide evidence of spillovers across contracts which are not managed by the robot, consistent with investors' learning.

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* .

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. (2021), 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.

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* .

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.

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.

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.

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.

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

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

Philippon, T. (2019), 'On fintech and financial inclusion', *NBER Working Paper No. 26330* .

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* .

7 Tables

Table 1: Descriptive Statistics

Variable	p5	mean	p95	sd	N
Age	30.00	49.24	67.00	11.74	5,878,593
Female	0.00	0.31	1.00	0.46	5,860,982
Saving plan value	0.00	7,541	36,074	26,387	5,879,049
Total account value	0.00	37,764	155,248	76,579	5,879,049
Yearly variable remuneration	0.00	1,748	7,832	3,263	5,879,049
Nb of saving vehicles	1.00	4.71	12.00	3.57	5,879,049
Number of connexions per month	0.00	0.85	4.00	3.13	2,263,612
Equity share	0.00	0.19	0.84	0.28	3,799,350
Number of asset allocation changes	0.00	0.05	0.00	0.25	5,879,049
Number of asset allocation changes (robo)	0.00	0.01	0.00	0.08	5,879,049
Number of asset allocation changes (indiv)	0.00	0.01	0.00	0.12	5,879,049
Annual return (realized)	-0.18	0.04	0.29	0.16	3,300,663
Annual return (expected)	-0.01	0.08	0.30	0.12	3,299,682
Volatility	0.00	0.14	0.44	0.67	3,300,663
Beta	0.00	0.95	2.89	2.11	2,827,366

NOTE: This table reports descriptive statistics of our variables. 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			Changes	Robo(>t)	Individual
Robo treated*after	0.299***	0.266**	0.138**	0.0904***	0.0437***	0.00542
	(0.0898)	(0.0984)	(0.0523)	(0.0161)	(0.00368)	(0.00364)
Sample		No rem	No Sub			
Observations	808,816	659,537	799,002	3,835,642	3,835,642	3,835,642
R-squared (within)	0.01	0.01	0.01	0.01	0.01	0.01

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 two months around 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. Columns 1-3 include individual and time fixed effects, columns 4-6 include contract and time fixed effects. 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, 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	Connexions	Rebalancer	Change in Distance Actual - Target Equity		
Robo treated*Alert	0.188***	0.284***	-0.0407***		
	(0.0599)	(0.0567)	(0.00236)		
Alert	0.298***	0.134***	0.0350***	-0.00613***	
	(0.0855)	(0.0154)	(0.00292)	(0.00144)	
Alert MIF					0.00124*
					(0.000721)
Sample	Robo takers+curious			Robo takers	
Observations	128,774	1,865,018	1,865,018	679,577	614,292
R-squared (within)	0.01	0.1	0.01	0.01	0.01

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. All regressions include time and contract fixed effects. 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, are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4: Returns

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Variable		Realized Return			Expected Return	
Robo treated*after	0.0421***	0.0371***	0.0433***	0.0486***	0.0285***	0.0215***
	(0.00789)	(0.00778)	(0.00791)	(0.00110)	(0.00446)	(0.00446)
Equity share		0.0892***				0.125***
		(0.00764)				(0.00687)
Volatility			0.0484***			
			(0.00224)			
Beta				0.0159***		
				(0.00155)		
Observations	3,279,169	3,279,169	3,279,169	2,804,489	3,275,769	3,275,769
R-squared (within)	0.01	0.01	0.06	0.05	0.01	0.03

NOTE: This table reports the results of OLS regressions. In columns 1-4, the dependent variable is the annual returns at the contract level. In columns 5-6, the dependent variable is the expected annual returns at the contract level. All regressions include contract and time fixed effects. 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, are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 5: Returns: Static vs. Dynamic Effect

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Variable	Static (Realized)				Static (Expected)	
Robo treated*after	0.0132***	0.00879***	0.0134***	0.0134***	0.0110***	0.00425**
	(0.000433)	(0.000415)	(0.000432)	(0.000458)	(0.00172)	(0.00192)
Equity share		0.0795***				0.122***
		(0.00128)				(0.0127)
Volatility			0.0101***			
			(0.000763)			
Beta				0.00275***		
				(0.000247)		
Observations	3,279,169	3,279,169	3,279,169	2,804,489	3,275,769	3,275,769
R-squared (within)	0.01	0.01	0.01	0.01	0.01	0.06

NOTE: This table reports the results of OLS regressions decomposing the total change in returns associated to the robo-service between a static effect occurring at the time of the subscription and a dynamic effect associated to different portfolio dynamics after the subscription. For robo-takers, define t* as the date of robo-subscription and, for non-exposed, as the date of first reception of the variable remuneration. In columns 1-4, the dependent variable is the static effect on annual returns, computed according to equation (4), which is the difference between the returns the investor has experienced after t* and those she would have experienced had she kept her portfolio weights constant at the level observed just before t*. In columns 5-6, the dependent variable is the same static effect, computed instead on expected returns. All regressions include contract and time fixed effects. 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, are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 6: Realized and Counterfactual Returns

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. Variable	Realized - Counterfactual Returns (in %)						
Robo rebalancing			0.229*		0.153**		0.120***
			(0.116)		(0.0721)		(0.0428)
Free rebalancing			-0.759		-0.434		0.0240
			(0.765)		(0.505)		(0.214)
Age		-0.0012					
		(0.0052)					
Female		0.0514					
		(0.0684)					
Account value (ln)		0.0142					
		(0.0223)					
Var. remuneration (k)		-0.0024*					
		(0.0013)					
Constant	-0.0447***	0.00592	-0.0438***	-0.0756***	-0.0761***	-0.0888***	-0.0939***
	(0.0101)	(0.238)	(0.0102)	(0.0137)	(0.00619)	(0.0139)	(0.00293)
Horizon		1-month		3-months		6-months	
Contract FE	No	No	Yes	No	Yes	No	Yes
Observations	642,048	618,756	642,048	606,291	606,291	552,735	552,735
R-squared (within)	0.01	0.01	0.01	0.01	0.01	0.01	0.01

NOTE: This table reports the results of OLS regressions in which the sample is restricted to robo-takers. The dependent variable is the difference in percentage points between the investor's annual returns and the counterfactual returns obtained by rebalancing the portfolio immediately upon reception of the alert as suggested by the robot. Robo rebalancing is a dummy equal to 1 if the investor rebalances as suggested by the robot, at any point in time. Free rebalancing is a dummy equal to 1 if the investor rebalances without following the robot's recommendation. In columns 1-3 we consider average returns in month following the rebalancing decision, in columns 4-5 we consider a 3-month horizon and in columns 6-7 a 6-month horizon. All regressions include time fixed effects. Standard errors, double-clustered by individual and time, 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		Trading		Returns		Static	
Robo treated*after	0.27**	0.29***	0.09***	0.09***	0.029***	0.029***	0.008***	0.007**
	(0.0974)	(0.0734)	(0.0181)	(0.0161)	(0.00469)	(0.00604)	(0.00174)	(0.00286)
Control group	Exposed	Curious	Exposed	Curious	Exposed	Curious	Exposed	Curious
Observations	831,846	879,065	3,676,837	4,768,327	3,184,405	4,129,632	3,184,405	4,129,632
R-squared (within)	0.01	0.01	0.01	0.01	0.01	0.01	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 dependents variable is the annual return; In columns 7 and 8, the dependent variable is the static effect as defined in equation (4). 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 expressed interest but did not take the robot (robo-curious). In columns 1-2, we include individual and time fixed effects; in columns 3-8, we include contract and time fixed effects. 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, are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 8: IV Estimates

	(1)	(2)	(3)	(4)
Dep. Variable	Connexions	Trading	Returns	Static
Robo treated*after	0.153*	0.0666**	0.0630***	0.0369***
	(0.0840)	(0.0267)	(0.0150)	(0.00561)
	First Stage: Robo Treated			
Fraction of treated employees	12.992***	3.225***	3.084**	3.084**
	(2.0707)	(0.3299)	(0.3617)	(0.3617)
F-Stat (first stage)	39.37	95.55	72.7	72.7
Observations	28,917	3,822,677	3,267,180	3,267,180
R-squared (within)	0.012	0.006	0.008	0.002

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 annual return; In column 4, the dependent variable is the static effect as defined in equation (4). In column 1, we include individual and time fixed effects; in columns 2-4, we include contract and time fixed effects. 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, are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.