Can Robot-Advising Reduce Inequalities?*

Milo Bianchi† Marie Brière‡

January 15, 2026

Abstract

We study how the introduction of a robot-advising service affects inequalities across investors. We find that the robot tends to equalize investment strategies and returns across investors, inducing larger increases in equity for those with lower baseline equity and larger increase in returns for those with lower baseline returns. At the same time, investors have different propensities to take-up the robot and to follow its recommendations. We show that younger, male and poorer investors experience larger increase in returns, and, for those investors, much of the effect is driven by changes in rebalancing strategies occurring over time. Access to the robot *decreases* the inequalities between poor and rich investors and between old and young investors, while at the same time *increasing* the inequalities associated to gender.

Keywords: Robot-Advising. Wealth Inequality

JEL codes: G11; G51; G41; G23; D14.

*Milo Bianchi acknowledges funding from 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

Households have different abilities to undertake complex financial decisions, leading to significant heterogeneity in the performance of their investments (Lusardi and Mitchell (2014), Campbell (2006)). Long-run patterns of wealth accumulation and inequality are strongly driven by the fact that investors with lower financial education earn significantly lower returns (Lusardi, Michaud and Mitchell (2017), Bianchi (2018)), while wealthier individuals earn persistently higher returns (Bach, Calvet and Sodini (2020), Fagereng, Guiso, Malacrino and Pistaferri (2020)). These differences are particularly important for women (Lusardi and Mitchell (2008), Bucher-Koenen, Alessie, Lusardi and van Rooij (2024)) and for older investors (Finke, Howe and Huston (2017), Mazzonna and Peracchi (2024)).

In this context, automated financial advisors, often called robot-advisors, have attracted 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)). This is particularly important for investors with lower access to traditional financial advice, lower experience, and lower ability to detect advisers' misconducts (Egan, Matvos and Seru (2019), Reuter and Schoar (2024)). In this way, robot-advisers may potentially temper the observed inequalities in returns. Their ultimate effects however depend on the extent to which robots can properly serve the needs of disadvantaged investors (Reher and Sokolinski (2024)) and on how much these investors are willing to rely on automated recommendations (Bhattacharya, Hackethal, Kaesler, Loos and Meyer (2012), Greig, Ramadorai, Rossi, Utkus and Walther (2022)).

We investigate these issues by exploiting the introduction of a robot-advising service in a large set of Employee Saving Plans. An important feature of our setting is that a large proportion of our investors have small portfolios and little experience in the stock market, making it an ideal sample to study whether access to the robot can significantly affect the behaviors of less sophisticated investors. Another key distinctive feature of this service is that it is truly a robot-advisor which gives advice to the investors, both at the time of the subscription and over time, while leaving investors free to follow or to ignore the advice. This makes it different from the more common robot-advisors that automate portfolio investment and rebalancing (as e.g. in Loos, Previtero, Scheurle and Hackethal (2020), Reher and Sokolinski (2024), Rossi and Utkus (2024)). These features allow investigating whether the robot may affect the patterns of financial inclusion and wealth inequalities even while letting investors retain control over their portfolios.

Whether robot-advising tempers or exacerbates existing inequalities is a key open question (Philippon (2019), Abraham, Schmukler and Tessada (2019)). The robot may in principle reduce inequalities by equalizing investment strategies across investors, in which case we would observe larger effects on investor with "less efficient" or "more extreme" allocations to start with (Rossi and Utkus (2024)). At the same time, while investors with lower financial capabilities may have the most to gain from robot-advising (D'Hondt, De Winne, Ghysels and Raymond (2020)), they may also be more reluctant to adopt the new technology (Hackethal, Haliassos and Jappelli (2012), Collins (2012), Foster and Rosenzweig (2010)), or they may end up misusing it (Campbell (2006)). Moreover, even if the robot may bring value "on average", it may be better suited to serve some types of investors than others or even give "biased" recommendations towards some groups, possibly resulting in increased inequalities across investors.

We mostly focus on inequalities among investors with different demographic characteristics, such as gender, age and wealth. This is motivated by the above evidence that differences in returns along these dimensions appear particularity relevant when analyzing the patterns of inequalities. At the same time, it is important to notice that the robot under study does not condition its recommendations on any of those characteristics. As we describe in more details below, the robot's recommendations depend only on investors' horizon and risk profile, irrespective of their demographics. Hence, from an ex-ante perspective, there is no mechanical reason why the robot would advise differently males and females, for example. However, ex-post, differential impacts may still arise if males and females receive different recommendations as they have different risk profiles, for example. Or, even if they get the same recommendations, differences may arise if they have different propensities to take-up the service or to rely on the robot over time. In this way, even if the robot's algorithm is "unaware" of demographic characteristics, it may still impact inequalities along these dimensions.

We start our analysis by observing significant heterogeneity in returns across investors before the introduction of the robot. These differences are correlated with demographic characteristics: poorer investors (as proxied by the value of their portfolio) and younger investors experience lower returns, while women do not display significant differences in returns relative to men. These patterns change after the introduction of the robot: the difference in returns between poorer and richer investors is reduced; younger investors experience larger returns than older ones; women experience lower returns than men. In our next analysis, we uncover the extent to which these patterns can be attributed to the robot and what are the underlying mechanisms.

One reason why investors may be impacted differently by the robot may simply be that they have different propensity to take-up the service. We show that indeed women are less likely to subscribe to the robot and to accept an increase in equity exposure: they are twice less likely to take-up when the robot proposes an increase in equity. We also notice that richer investors show lower take-up. Other characteristics have minimal effects.

Another reason is that investors receive different risk recommendations from the robot or that they have different propensities to change their risk exposure following the robot. We show that the robot tends to equalize equity exposure across investors, inducing a larger increase for those with lower equity at the baseline. Investors with 0% equity at the baseline (corresponding to the 25th percentile in the distribution of equity exposure) increase their equity by 13.5% while those with 22.5% equity (the 75th percentile) increase it by 1.3%. Significant heterogeneity is also observed along demographic characteristics: the largest increase in equity is observed for men, younger, and poorer investors.

We show that much of these changes in equity occur at the take-up period, as result of the robot's recommendations. These are by construction determined according to two dimensions: investment horizon, as reported by each investor, and a risk index, a composite measure of self-reported risk tolerance and financial knowledge. The robot proposes lower equity to older investors, as they have shorter horizon and lower risk tolerance, despite that at the baseline these investors hold more equity than younger investors. The robot proposes a larger increase in equity to poorer investors, thereby flattening the relation between equity exposure and wealth. The robot proposes lower equity to females as they are estimated to have a lower tolerance towards risk, even if at the baseline they display no significant difference with male investors.

These changes in investment strategies have significant impacts on returns. The robot has an equalizing effects on returns: robo-takers with high baseline returns (equal to 8.7%, corresponding to the 75th percentile) experience a 1.17% increase in returns, while those with low baseline returns (equal to 0.2%, corresponding to the 25th percentile) experience a 5.95% increase. That is, the inequality in returns is reduced to less than half of its baseline value. The effects related to the other characteristics follow similar patterns as those on equity. Female investors experience a lower increase in returns than males; younger investors experience a higher increase in returns than older investors; poorer investors experience a higher increase than richer investors.

A third reason behind these heterogeneous effects is that, after take-up, investors have a different propensity to change the way in which they rebalance their portfolios. We decompose the total effect on returns 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. In order to isolate the dynamic effect, we identify a subset of robo-curious investors whose portfolio allocation when completing the profiling survey is close to the one suggested by the robot. For these investors, the potential changes induced by the robot would have been essentially those associated with rebalancing behaviors after subscription. Using these investors as reference point, we show that for those investors who experience the largest increase in returns (that is, male, younger and poorer investors) the dynamic effect is the most important driver of this increase.

These effects have important implications first as they portray a rich picture in terms of inequalities. Access to the robot significantly *decreases* the inequalities between poor and rich investors and between old and young investors, while at the same time *increasing* the inequalities associated to gender. This is the case even if, as mentioned, the robot does not explicitly condition on these characteristics, showing how even ex-ante "neutral" algorithms may affect inequalities.

In the concluding section, we assess these effect from a quantitative viewpoint, combining both the different individual propensities to take-up the robot and the observed changes in returns conditional on take-up. We show that the robot is essentially creating gender inequalities: before take-up, female investors had basically the same (if anything, slightly larger) returns than males, but after the introduction of the robot male investors experience 0.48% larger returns. At the same time, the introduction of the robot is associated with a 74% decrease in inequality in returns between old and young investors and to a 37% decrease in inequality in returns between rich and poor investors.

Our results are also important for studying financial markets more broadly. Our estimates suggest that access to the robot, if generalized, may affect the dynamics of asset prices by changing the composition of equity holding across various groups of investors. Recent evidence shows for example that financial markets react differently to macro shocks depending on the composition of investors in terms of gender (Elkind, Kaminski, Lo, Siah and Wong (2022)), age (Malmendier, Pouzo and Vanasco (2020)), and wealth (Hoopes, Langetieg, Nagel, Reck, Slemrod and Stuart (2022)).

Finally, our insights may inform the policy debate on the extent to which financial institutions and FinTechs should be required to test the implications of their algorithms and avoid amplifying existing inequalities. As we discuss in the concluding remarks, these results are also an important building block when assessing the optimal personalization of the algorithms' recommendations.

2 Data

2.1 Investment Plan and Robo-Advisor

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.[1]

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.[2]

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.[3]

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

[2] 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.

[3] Even if the investor accepts the robot's allocation, she is not committed to it in any way, she can

Over time, the robot also sends email alerts if the current portfolio allocation ends up being too far from the target allocation.

2.2 Sample

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, based on the information on the 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' trading activities on the platform. 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.[4] We provide more details about these 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 curious. We restrict to individuals who have completed at least one transaction in one of their 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.

In Table 1, we report some descriptive statistics on demographic and portfolio characteristics of our sample of investors. For baseline characteristics, we aggregate variables at the individual level and, for each investor, we consider the average value of the variable

change again the allocation right after having taken up the robot.

[4]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.

before the introduction of the robo-service, between January 2016 and August 2017. 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.[5] 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 Inequalities Before and After the Robot

We start by providing some descriptive evidence that motivates our next analysis. While in our setting investors allocate their wealth between a menu of funds, which relative to stock picking should minimize issues of under-performance and under-diversification, we observe very large heterogeneity across investors' returns before the introduction of the robot. In our sample, the average return is 3.5% and the standard deviation is almost four times larger, about 13.2%. Moreover, about one third of this standard deviation is between investors, indicating that at any given point in time, some investors experience significantly larger returns than others.

These differences are systematically correlated with equity exposure and with demographic characteristics. Moreover, remarkably, the relations between demographic characteristics, equity and returns change with the introduction of the robot. In Figure 1, we plot the average equity exposure in our sample (including both robo-curious and robo-takers) across various groups of investors before and after the introduction of the robot, in August 2017. Equity before refers to the average equity exposure over the 10 months in our sample before the robot, from September 2016 and July 2017. Equity after is instead computed as the average equity exposure over the 10 months at the end of our sample, from December 2019 and June 2021. We first observe that the distribution of equity exposure is much more dispersed before than after the robot. Investors in the first and second quartile of the baseline equity distribution are those who experience the largest increase in equity after the robot; while those in the last quartile decrease their

[5] 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).

equity exposure.

In terms of demographics, we notice that the relation between equity exposure and age is increasing before the robot, while it is much flatter after the robot. There is no significant increase in equity for the oldest quartile (investors aged 58 and older), in contrast with younger investors. We also notice that while there is no significant difference in equity exposure between men and women before the robot, men display larger equity after the robot. Finally, the relation between investors' wealth (as proxied by the value of their portfolio) and equity is increasing, but much flatter after the robot.

In Figure 2, we plot the same figures in terms of annual returns. As for equity, the distribution of returns is much less dispersed after the robot. Investors in the lower quartile in the baseline distribution experience the largest increase in returns after the robot, investors in the highest quartile experience a decrease in returns. Similar patterns are observed in terms of demographics. Returns increase with age before the robot, while they decrease with age after the robot. Returns do not significantly depend on gender before the robot, while they are significantly higher for males after the robot. Wealthier investors tend to experience higher returns, but the relation is much weaker after the robot.

These figures suggest that the patterns of inequalities display significant differences after the introduction of the robot. It should be however noticed that these figures report raw data on the entire sample of investors, without accounting for actual take-up of the robot, time-trends or other confounding factors. In our next analysis, we uncover the extent to which these patterns can be attributed to the robot and what are the underlying mechanisms.

4 Accessing the Robot

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

$$y_{i,t} = \alpha_i + \beta T_{i,t} + \gamma T_{i,t} X_i + \mu_t + \varepsilon_{i,t}, \tag{1}$$

where α_i 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.

$T_{i,t}$ is a dummy equal to 1 if the investor has taken the robot in contract i at time t and X_i are characteristics, taken as averages at the baseline (September 2016-July 2017) : gender, equity exposure, returns, age, wealth (portfolio value). We double cluster standard errors by individual and time.

To facilitate the interpretation of our estimates, unless specified otherwise, we will compare investors sitting at the 25th percentile of the distribution to to those sitting at the 75th percentile. In terms of age, we compare 54 years old investors, labeled "old", to 36 years old ones, labeled "young"; in terms of total portfolio value, we compare those with 845 euros (6.74 in log terms), labeled as "poor", to those with 16,800 euros (9.73 in log terms), labeled as "rich". Similarly, in terms of baseline equity, we compare those with 0% to those with 22.5%; in terms of baseline returns, we compare investors with 0.2% to those with 8.7%.

4.1 Take-Up

In Table 2, we investigate whether the propensity to take-up the robot is associated with our characteristics of interest by comparing those who have subscribed to the service (robo-takers) to those who have initiated the profiling process without eventually subscribing (robo-curious).[6] Older, female and richer investors are less likely to take-up the robot. In terms of magnitudes, the largest differences are driven by gender and wealth. About 42% of male investors have taken up the robot in this sample, compared to 37% of female investors. For rich investors, the propensity to take-up is 35% while for poor investors it is 45%. Other characteristics have smaller effects. Young investors have a 41% take-up rate, compared to 39% for old investors. The effects of ex-ante equity and ex-ante returns are also small: a standard deviation increase in those variable is associated to a decrease in take-up around 1%, even if, as shown in the previous section, investors with lower equity and lower returns would have potentially more to gain from accessing the robot.

In column 6, we investigate how the robot's equity recommendation affects the decision to take-up. We interact our variables of interests with a dummy equal to one if the robot proposes an increase in equity exposure. We observe that older, female, poorer investors are more reluctant to take-up the robot when it proposes an increase in equity exposure; while investors with higher baseline equity exposure and returns are more likely to accept such equity increase. In terms of magnitudes, we notice a large effect on female, who are

[6]While for simplicity we consider linear regressions here, probit specifications give very similar results.

twice less likely to take-up when offered an increase in equity.

4.2 Equity Exposure

While the adoption of the robot tends to be associated with an increase in the exposure to equity (Bianchi and Brière (2024)), we observe a large heterogeneity across investors. The robot tends to equalize the equity exposure, inducing larger changes for those investors with lower equity at the baseline. Specifically, as shown in Table 3, those with 0% equity at the baseline (corresponding to the 25th percentile in the distribution of equity exposure) increase their equity by 13.5% while those with 22.5% equity (the 75th percentile) increase it by 1.3%.

Significant heterogeneity is also observed along other dimensions. For men, the increase is 5.4% while for women it is 2.4%; for rich investors, the increase is 2.6%, while for poor investors it is 6.8%; for the old it is 0.7% while for the young it is 7.9%. That is, men, poorer and younger investors display larger increase in equity exposure upon adoption of the robo-advice.[7]

We then investigate how much of these changes is driven by the robot's recommendation at take-up, as opposed to changes occurring after subscription (and possibly unrelated to the robot). In Table 4, we restrict the sample to all periods up and including the take-up period t^*. That is, while in Table 3 we compare the average equity exposure in all periods $t < t^*$ with all periods $t \geq t^*$, we here omit all periods $t > t^*$ and compare the average equity exposure before t^* to the one at t^*. We observe that the coefficients in Table 4 are almost identical to the corresponding ones in Table 3, showing that the observed changes in equity exposure essentially occur at take-up.

As much of the effect is driven by investors' behavior at take-up, and so by the robot's equity recommendations, we now analyze the determinant of these recommendations. By design, they depend on investors' reported horizon, measured in years, and by their risk profile, which is based on reported risk tolerance and financial knowledge and measured on a 1-5 index. In our sample, the standard deviation of investors' horizon is 6 years, while for the risk index it is 1.15. Comparing investors with high (22.5%) vs. low (0%) equity exposure at the baseline, we observe in Table 5 that those with higher equity have a 0.68 longer horizon and a 0.26 larger risk profile, which translates into a 5.3% larger equity recommendation. This is however much smaller than the differences between the

[7] We find similar effects when considering the probability to hold equity ("stock market participation"), rather than the continuous measure of equity exposure.

two groups before take-up; in this way, the robot tends to compress the distribution of equity exposure.

Similar dynamics occur relative to demographic characteristics. The robot proposes lower equity to older investors, as they have shorter horizon and lower risk tolerance, even if at the baseline older investors tend to hold more equity than younger ones. The robot proposes higher levels of equity to richer investors, but the proposed increase in equity is larger for poorer investors, thereby tempering the differences in equity exposure relative to wealth. The robot proposes lower equity to females, essentially as they are estimated to have a lower tolerance towards risk, even if at the baseline they display no significant difference with male investors.[8]

4.3 Returns

We now investigate how the changes in investment strategies described above translate into changes in returns. We consider annual returns net of fees. Bianchi and Brière (2024) show that on average robo-takers experience larger returns; in Table 6, we show that these effects are highly heterogeneous across investors. We start by comparing investors with high baseline returns (8.7%, corresponding to the 75th percentile) to those with low baseline returns (0.2%, corresponding to the 25th percentile). The robot has an equalizing effects on returns: robo-takers with high baseline returns experience a 1.17% increase in returns, while those with low baseline returns experience a 5.95% increase. That is, the inequality in returns is reduced to less than half of its baseline value.

The effects related to the other characteristics follow similar patterns as those related to equity. Female investors experience a lower increase in returns, equal to 1.97%, as compared to 2.86% for males. Young investors experience a higher increase in returns, equal to 3.31%, as compared to 2.08% for old investors. Poor investors also experience a higher increase in returns, equal to 3.45%, as compared to 2.03% for rich investors.

These effects may be in principle driven by two different changes in investors' strategies. 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 the next analysis, 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

[8]The low risk index for female may be related to their low confidence in financial literacy questions, as documented in Bucher-Koenen et al. (2024).

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.

4.3.1 Static and Dynamic Effects

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 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.

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.[9] We repeat the analysis of Table 6 while using the sample of curious close as control; results are reported in Table 7.

We observe that the relative importance of the dynamic effect is again highly heterogeneous and, remarkably, it is significantly higher for those investors who experience the largest increase in returns after having taken up the robot. Specifically, the dynamic effect accounts for about 12% of the total effect for investors with high baseline returns,

[9] We also perform robustness checks using a 10% threshold in the distance, corresponding to 15% of the robo-curious.

while it accounts for 76% of the total effect for investors with low baseline returns.[10] That is, the increase in returns experienced by those investors is mostly driven by changes in the way they rebalance their portfolio over time. A similar result is observed in Bianchi (2018) for investors with lower financial literacy.

Similar patterns emerge for the other groups of investors. For young investors, the dynamic effect accounts for 58% of the total effect, compared to 32% for old investors. For males, the dynamic effect accounts for 52% of the total effect, compared to 31% for females. For poor investors, the dynamic effect accounts for 60% of the total effect, compared to 28% for rich investors. Overall, the dynamic effect is the most important driver of the increase in returns for those investors who experience the largest increase.

5 Discussion

We have shown that accessing the robot has significantly different effects across investors. In general, the robot tend to equalize investment strategies and returns across investors, inducing larger increases in equity for those with lower baseline equity and larger increase in returns for those with lower baseline returns. At the same time, investors have different propensity to follow the robot. We show that younger, male and poorer investors experience larger increase in returns, and, for those investors, much of the effect is driven by changes in rebalancing strategies occurring after the robot take-up. Access to the robot *decreases* the inequalities between poor and rich investors and between old and young investors while at the same time *increasing* the inequalities associated to gender.

In order to get a sense of the magnitudes, we report in Table 8 how the inequalities in returns change before and after the introduction of the robot. Relative to Figure 2, we now consider returns and take-up at the individual level, and we evaluate the effects of the robot building on our regressions, combining both the probability of take-up and the effects conditional on take-up. We observe that before take-up female investors had essentially the same (if anything, slightly larger) returns than males, but after the robot the expected change in returns is 0.48% larger for male investors, leading to 0.43% larger returns for males after take-up. In this way, the robot is essentially creating gender inequalities in returns, and the effect is driven by the fact that male investors have both a higher probability of take-up and a larger increase in returns conditional on take-up.

[10] For investors with low baseline returns, the dynamic effect is 4.55% (that is, $4.66 - 55.96 * 0.002$) while the total effect is 5.95%, hence the dynamic effect corresponds to 76% of the total. We perform similar computations for the other groups of investors.

Older investors experience 0.74% larger returns before take-up, and 0.55% lower returns after the introduction of the robot, resulting in 0.19% higher returns after take-up. This corresponds to a 74% decrease in inequality in returns relative to the baseline, and the effect is driven by a slightly lower probability of take-up and especially by a lower increase in returns conditional on take-up for older investors. Finally, richer investors experience 2.27% larger returns before take-up, and 0.84% *lower* returns after the introduction of the robot, leading to 1.43% higher returns after take-up. This corresponds to a 37% decrease in inequality in returns relative to the baseline, which results from a lower probability of take-up and a lower increase in returns conditional on take-up for richer investors.

We view our analysis only as a first step in the investigation of how automated financial services may impact inequalities. An important element in this discussion is the extent to which algorithms which are ex-ante fair, as they treat all investors on equal ground, end up having distortionary effects and they may exacerbate or even create inequalities. For example, in our setting, the robot does not condition its recommendations on gender, but it ends up recommending lower equity to women as they report a lower capacity to bear risk. At the same time, women tend to be less confident, which may affect their self-reported risk ability. An open question is whether providing recommendations tailored on gender would affect the take-up and the willingness to rely on the robot. More generally, addressing the question of which characteristics the algorithm should take into account, and so what is the optimal degree of personalization, requires assessing how personalization may affect different groups in the population and ultimately inequalities.

References

Abraham, F., Schmukler, S. L. and Tessada, J. (2019), 'Robo-advisors: Investing through machines', *World Bank Research and Policy Briefs* (134881).

Bach, L., Calvet, L. E. and Sodini, P. (2020), 'Rich pickings? risk, return, and skill in household wealth', *American Economic Review* **110**(9), 2703–47.

Bhattacharya, U., Hackethal, A., Kaesler, S., Loos, B. and Meyer, S. (2012), 'Is unbiased financial advice to retail investors sufficient? answers from a large field study', *Review of Financial Studies* **25**(4), 975–1032.

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.

Bianchi, M. and Brière, M. (2024), 'Human-robot interactions in investment decisions', *Management Science* .

Bucher-Koenen, T., Alessie, R., Lusardi, A. and van Rooij, M. (2024), 'Fearless woman: Financial literacy, confidence, and stock market participation', *Management Science, forthcoming* .

Campbell, J. Y. (2006), 'Household finance', *Journal of Finance* **61**(4), 1553–1604.

Collins, J. M. (2012), 'Financial advice: A substitute for financial literacy?', *Financial Services Review* **21**(4), 307.

D'Acunto, F. and Rossi, A. G. (2023), 'Robo-advice: Transforming households into rational economic agents', *Annual Review of Financial Economics* **15**(1), 543–563.

D'Hondt, C., De Winne, R., Ghysels, E. and Raymond, S. (2020), 'Artificial intelligence alter egos: Who might benefit from robo-investing?', *Journal of Empirical Finance* **59**, 278–299.

Egan, M., Matvos, G. and Seru, A. (2019), 'The market for financial adviser misconduct', *Journal of Political Economy* **127**(1), 233–295.

Elkind, D., Kaminski, K., Lo, A. W., Siah, K. W. and Wong, C. H. (2022), 'When do investors freak out? machine learning predictions of panic selling', *Available at SSRN 3898940* .

Fagereng, A., Guiso, L., Malacrino, D. and Pistaferri, L. (2020), 'Heterogeneity and persistence in returns to wealth', *Econometrica* **88**(1), 115–170.

Finke, M. S., Howe, J. S. and Huston, S. J. (2017), 'Old age and the decline in financial literacy', *Management Science* **63**(1), 213–230.

Foster, A. D. and Rosenzweig, M. R. (2010), 'Microeconomics of technology adoption', *Annual Review of Economics* **2**(1), 395–424.

Greig, F., Ramadorai, T., Rossi, A. G., Utkus, S. P. and Walther, A. (2022), 'Algorithm aversion: Theory and evidence from robo-advice', *Available at SSRN 4301514* .

Hackethal, A., Haliassos, M. and Jappelli, T. (2012), 'Financial advisors: A case of babysitters?', *Journal of Banking & Finance* **36**(2), 509–524.

Hoopes, J. L., Langetieg, P., Nagel, S., Reck, D., Slemrod, J. and Stuart, B. A. (2022), 'Who sells during a crash? evidence from tax return data on daily sales of stock', *The Economic Journal* **132**(641), 299–325.

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., Michaud, P.-C. and Mitchell, O. S. (2017), 'Optimal financial knowledge and wealth inequality', *Journal of Political Economy* **125**(2), 431–477.

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

Lusardi, A. and Mitchell, O. S. (2014), 'The economic importance of financial literacy: Theory and evidence', *Journal of Economic Literature* **52**(1), 5–44.

Malmendier, U., Pouzo, D. and Vanasco, V. (2020), 'Investor experiences and financial market dynamics', *Journal of Financial Economics* **136**(3), 597–622.

Mazzonna, F. and Peracchi, F. (2024), 'Are older people aware of their cognitive decline? misperception and financial decision-making', *Journal of Political Economy* **132**(6), 1793–1830.

Philippon, T. (2019), *The great reversal: how America gave up on free markets*, Harvard University Press.

Reher, M. and Sokolinski, S. (2024), 'Robo advisors and access to wealth management', *Journal of Financial Economics* **155**, 103829.

Reuter, J. and Schoar, A. (2024), 'Demand-side and supply-side constraints in the market for financial advice'.

Rossi, A. G. and Utkus, S. (2024), 'The diversification and welfare effects of robo-advising', *Journal of Financial Economics* **157**, 103869.

Figures and Tables

Figure 1: Equity Exposure Before and After the Robot

NOTE: This figure displays the equity exposure in our sample (including both robo-curious and robo-takers) across various groups of investors before and after the introduction of the robot, in August 2017. Equity before is the average equity exposure between September 2016 and July 2017; Equity after is the average equity exposure between December 2019 and June 2021. The sample includes both robo-takers and robo-curious.

Figure 2: Returns Before and After the Robot

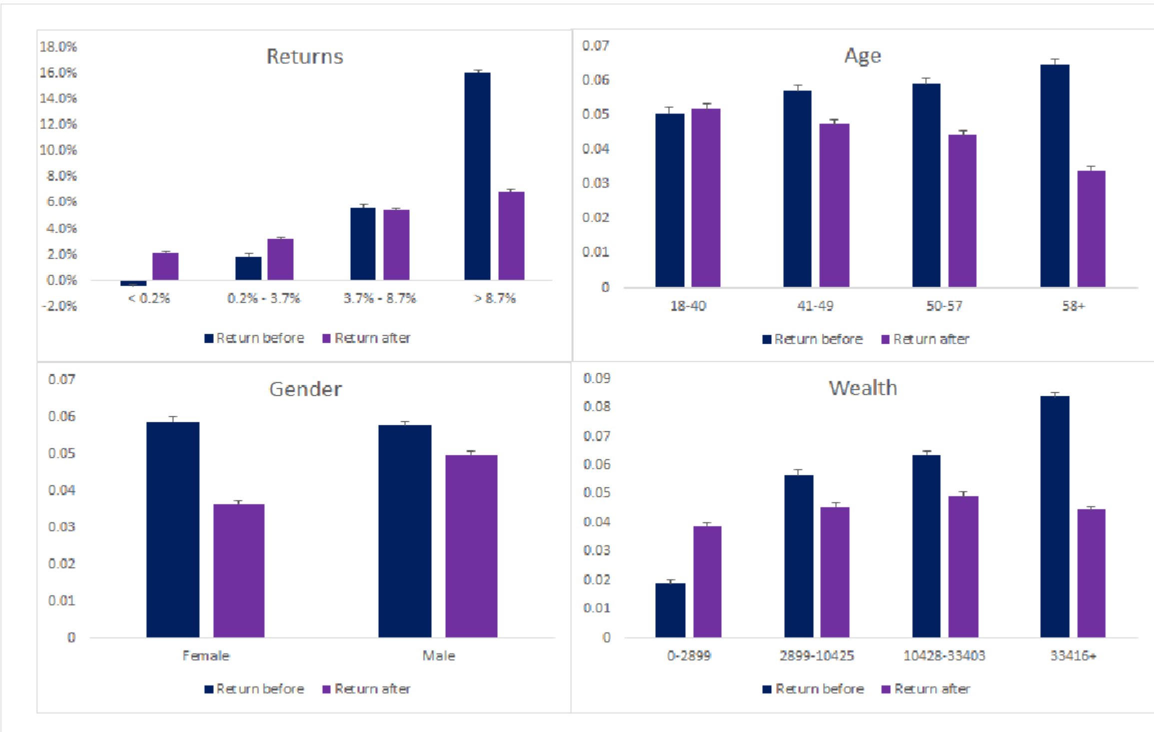

Note: This figure displays the annual returns in our sample (including both robo-curious and robo-takers) across various groups of investors before and after the introduction of the robot, in August 2017. Return before is the average return between September 2016 and July 2017; Return after is the average return between December 2019 and June 2021. The sample includes both robo-takers and robo-curious.

Table 1: Descriptive Statistics

All sample	p5	Mean	p95	SD	N
Age	30.00	48.362	65.00	11.18	4795438
Female	0.00	0.362	1.00	0.48	4795438
Total account value	147	34642	138967	61458	4710383
Total account value (ln)	6.15	9.330	11.86	1.79	4560457
Yearly variable remuneration	0.00	2405	9270	3690	4795438
Equity share	0.00	0.216	0.80	0.26	2782081
Proposed equity share	0.03	0.250	0.66	0.20	112878
Horizon	1.00	6.797	20.00	6.01	113645
Risk profile	1.00	2.264	4.00	1.15	113645
Proposed change in equity exposure	-0.32	0.014	0.36	0.20	111556
Ann. return	-0.09	0.035	0.24	0.13	3174911

At baseline	p5	Mean	p95	SD	N
Age	30.00	48.439	65.00	10.98	1074873
Female	0.00	0.375	1.00	0.48	1074873
Total account value	340	30364	115385	50936	1003830
Total account value (ln)	6.10	9.238	11.66	1.74	990863
Yearly variable remuneration	0.00	2452	8887	3478	1074873
Equity share	0.00	0.179	0.78	0.25	272803
Ann. return	-0.01	0.069	0.35	0.14	502153

NOTE: This table reports descriptive statistics of our variables. In Panel A, we report statistics all the observations by contract and over time for robo-takers and robo-curious, which are used in our main analysis. In Panel B, we restrict to periods before the introduction of the robot (September 2016 – July 2017).

Table 2: Take-Up

	(1)	(2)	(3)	(4)	(5)	(6)
			Robo take-up			
Age	-0.001***			0.001***	0.001***	0.003***
	(0.000)			(0.000)	(0.000)	(0.000)
Female		-0.053***		-0.050***	-0.048***	-0.021**
		(0.005)		(0.005)	(0.005)	(0.010)
Wealth			-0.030***	-0.035***	-0.028***	-0.042***
			(0.001)	(0.001)	(0.002)	(0.003)
Equity				0.049***	0.041***	-0.026
				(0.012)	(0.013)	(0.021)
Return					0.043	-0.388***
					(0.031)	(0.070)
Age*Increase						-0.001**
						(0.001)
Female*Increase						-0.024**
						(0.012)
Wealth*Increase						0.015***
						(0.004)
Equity*Increase						0.104***
						(0.029)
Return*Increase						0.184**
						(0.079)
Increase						0.005
						(0.030)
Constant	0.451***	0.419***	0.652***	0.636***	0.572***	0.882***
	(0.009)	(0.003)	(0.011)	(0.012)	(0.013)	(0.024)
Observations	37,149	38,014	36,823	36,823	33,485	28,420
R-squared	0.01	0.01	0.02	0.02	0.01	0.03

NOTE: The dependent variable is a dummy equal to 1 if the individual is a robo-taker and to zero if the individual is a robo-curious. All variables are computed as average at baseline, before the introduction of the robot (09/2016 – 07/2017). Wealth refers to the average value of the portfolio (in log). Increase is a dummy equal to one if the robot proposes an increase in equity share. Standard errors are in parenthesis. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 3: Change in Equity Exposure

	(1)	(2)	(3)	(4)	(5)	(6)
			Equity Exposure			
Robot*Equity	-0.543***					-0.543***
	(0.013)					(0.013)
Robot*Age		-0.004***			-0.004***	-0.004***
		(0.000)			(0.000)	(0.000)
Robot*Female			-0.030***		-0.025***	-0.035***
			(0.004)		(0.004)	(0.004)
Robot*Wealth				-0.014***	-0.004***	0.005***
				(0.001)	(0.001)	(0.001)
Robot	0.135***	0.223***	0.054***	0.162***	0.251***	0.265***
	(0.003)	(0.009)	(0.004)	(0.011)	(0.012)	(0.010)
Constant	0.206***	0.206***	0.207***	0.206***	0.206***	0.206***
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)
Contract FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,553,687	3,553,687	3,603,429	3,535,921	3,535,921	3,535,921
R-squared	0.08	0.02	0.01	0.01	0.02	0.09

NOTE: This table reports the results of OLS regressions. The dependent variable is the value of equity over the total value of the portfolio. Robot is a dummy equal to 1 if the investor has taken up the robot. Equity refers to the average value of equity over the total value of the portfolio at baseline, before the introduction of the robot (September 2016 – July 2017). Age refers to the average age of the investors at baseline. Wealth refers to the average value of the portfolio (in log) at baseline. 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: Change Equity Exposure at take-up

	(1)	(2)	(3)	(4)	(5)	(6)
			Equity Exposure			
Robot*Equity	-0.513***					-0.511***
	(0.016)					(0.014)
Robot*Age		-0.003***			-0.002***	-0.002***
		(0.000)			(0.000)	(0.000)
Robot*Female			-0.034***		-0.030***	-0.039***
			(0.009)		(0.008)	(0.006)
Robot*Wealth				-0.015***	-0.008	0.001
				(0.004)	(0.005)	(0.003)
Robot	0.130***	0.181***	0.056***	0.168***	0.229***	0.241***
	(0.010)	(0.016)	(0.015)	(0.045)	(0.040)	(0.026)
Constant	0.208***	0.208***	0.209***	0.208***	0.208***	0.208***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Contract FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,866,938	2,866,938	2,898,190	2,855,498	2,855,498	2,855,498
R-squared	0.01	0.01	0.01	0.01	0.01	0.01

NOTE: This table reports the results of OLS regressions. The dependent variable is the value of equity over the total value of the portfolio and the sample includes all periods up and including the take-up period. Robot is a dummy equal to 1 if the investor has taken up the robot. Equity refers to the average value of equity over the total value of the portfolio at baseline, before the introduction of the robot (September 2016 – July 2017). Age refers to the average age of the investors at baseline. Wealth refers to the average value of the portfolio (in log) at baseline. 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: Proposed Equity Exposure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Proposed equity share					Horizon	Risk
Equity	0.238***				0.215***	3.031***	1.146***
	(0.009)				(0.008)	(0.260)	(0.045)
Age		-0.001***			-0.003***	-0.091***	-0.019***
		(0.000)			(0.000)	(0.004)	(0.001)
Female			-0.057***		-0.052***	-0.543***	-0.402***
			(0.003)		(0.003)	(0.101)	(0.017)
Wealth				0.018***	0.022***	0.475***	0.068***
				(0.001)	(0.001)	(0.029)	(0.005)
Constant	0.223***	0.297***	0.274***	0.111***	0.178***	6.947***	2.544***
	(0.002)	(0.006)	(0.002)	(0.007)	(0.007)	(0.242)	(0.041)
Observations	18,227	18,227	18,752	18,029	18,029	18,042	18,042
R-squared	0.05	0.00	0.02	0.03	0.10	0.04	0.10

Note: This table reports the results of OLS regressions. In columns 1-5, the dependent variable is the equity share proposed by the robot at take-up. In column 6, the dependent variable is the investment horizon (in years) that the investor reports to the robot. In column 7, the dependent variable is the investor's risk profile, built by the robot based on the investor's reported risk tolerance and financial knowledge. Equity refers to the average value of equity over the total value of the portfolio at baseline, before the introduction of the robot (September 2016 – July 2017). Age refers to the average age of the investors at baseline. Wealth refers to the average value of the portfolio (in log) at baseline. 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: Change in Returns

	(1)	(2)	(3)	(4)	(5)	(6)
	Annual Returns					
Robot*Returns	-56.225***					-57.412***
	(7.181)					(7.145)
Robot*Age		-0.068***			-0.040***	-0.072***
		(0.013)			(0.014)	(0.013)
Robot*Female			-0.895***		-0.812***	-0.803***
			(0.294)		(0.286)	(0.262)
Robot*Wealth				-0.476***	-0.359***	0.284***
				(0.118)	(0.130)	(0.074)
Robot	6.066***	5.756***	2.862***	6.657***	7.777***	7.324***
	(0.865)	(0.877)	(0.459)	(1.342)	(1.330)	(1.194)
Constant	2.903***	2.851***	2.893***	2.837***	2.833***	2.902***
	(0.104)	(0.109)	(0.109)	(0.112)	(0.112)	(0.103)
Contract FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,962,597	3,134,111	3,174,911	3,119,296	3,119,296	2,962,597
R-squared	0.01	0.01	0.01	0.01	0.01	0.01

NOTE: This table reports the results of OLS regressions. The dependent variable is the annual returns at the contract level. Robot is a dummy equal to 1 if the investor has taken up the robot. Equity refers to the average value of equity over the total value of the portfolio at baseline, before the introduction of the robot (September 2016 – July 2017). Age refers to the average age of the investors at baseline. Wealth refers to the average value of the portfolio (in log) at baseline. 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: Returns: Comparing to Curious-close Investors

	(1)	(2)	(3)	(4)	(5)	(6)
			Annual Returns			
Robot*Returns	-55.960***					-57.064***
	(7.162)					(7.154)
Robot*Age		-0.070***			-0.039***	-0.069***
		(0.013)			(0.014)	(0.012)
Robot*Female			-0.878***		-0.801***	-0.800***
			(0.281)		(0.274)	(0.250)
Robot*Wealth				-0.510***	-0.396***	0.267***
				(0.108)	(0.119)	(0.065)
Robot	4.661***	4.450***	1.497***	5.529***	6.620***	5.904***
	(0.673)	(0.801)	(0.312)	(1.141)	(1.170)	(1.035)
Constant	2.459***	2.477***	2.548***	2.445***	2.438***	2.455***
	(0.171)	(0.175)	(0.173)	(0.182)	(0.181)	(0.172)
Contract FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,143,195	1,253,971	1,275,225	1,245,967	1,245,967	1,143,195
R-squared	0.04	0.01	0.01	0.01	0.01	0.04

NOTE: This table reports the results of OLS regressions. The dependent variable is the annual returns at the contract level. 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%. Robot is a dummy equal to 1 if the investor has taken up the robot. Equity refers to the average value of equity over the total value of the portfolio at baseline, before the introduction of the robot (September 2016 – July 2017). Age refers to the average age of the investors at baseline. Wealth refers to the average value of the portfolio (in log) at baseline. 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 8: Inequalities Before and After the Robot

	Female	Male	Difference
Returns (before take-up)	1.90	1.86	-0.04
Prob(take-up)	0.37	0.42	0.05
Change in Returns	1.97	2.86	0.89
Expected Change in Returns	0.72	1.20	0.48
Returns (after take-up)	2.63	3.06	0.43

	Young	Old	Difference
Returns (before take-up)	1.51	2.25	0.74
Prob(take-up)	0.41	0.39	-0.02
Change in Returns	3.31	2.08	-1.23
Expected Change in Returns	1.36	0.81	-0.55
Returns (after take-up)	2.87	3.06	0.19

	Poor	Rich	Difference
Returns (before take-up)	0.74	3.01	2.27
Prob(take-up)	0.45	0.35	-0.10
Change in Returns	3.45	2.03	-1.42
Expected Change in Returns	1.55	0.71	-0.84
Returns (after take-up)	2.29	3.72	1.43

NOTE: In this table, we report inequalities in returns before and after the introduction of the robot. Returns (before take-up) considers the average return, for each individual, over all periods before the robot take-up. Prob(take-up) in the probability of take-up, estimated in Table 2. Change in Returns is the impact the robot, as estimated in Table 6. Expected Change in returns is the product between Prob(take-up) and Change in Returns. Returns (after take-up) is the sum of Returns (before take-up) and Expected Change in Returns.