On the Benefits of Robo-Advice in Financial Markets Benefits of Robo-Advice

Accepted for publication in the Economic Journal

http://doi.org/10.1093/ej/ueag076

On the Benefits of Robo-Advice in Financial Markets

Benefits of Robo-Advice

Marco Lambrecht*, Joerg Oechssler, and Simon Weidenholzer

April 30, 2026

Abstract

Robo-advisors are tools in financial markets that provide investors with low-cost financial advice, typically based on individual characteristics such as risk attitudes. We study the benefits of robo-advice in a ten-week portfolio choice experiment. Depending on treatment, investors either receive robo-advice, have a robo-advisor implement recommendations by default, or invest on their own. While we observe no effect of robo-advice on initial market participation, we find positive effects on continued participation. Robo-advisors also help investors avoid mistakes, increase rebalancing, and yield portfolios closer to the utility-maximising benchmark. Default implementation of recommendations performs significantly better than advice alone.

JEL codes: C91, D81, G12, G20, G41.

Keywords: algorithmic trading, experiment, financial markets.

* *Corresponding author*: Department of Economics & Finance, University of Stavanger, PB 8002, 4036 Stavanger, Norway, email: marco.lambrecht@uis.no

Acknowledgements: We would like to thank seminar audiences at Aalto Finance, Ben Gurion University, Bern, Durham, Heidelberg, Helsinki GSE, Nottingham, Otago, WU Vienna, ESA Bologna, SEET Valencia, and the Experimental Finance Conference in Sofia for comments. We are also thankful to Mikhail Anufriev, Jose Apesteguia, Sebastian Ebert, Yohanes Eko Riyanto, Jean Paul Rabanal, and Martin Weber for very useful discussions. Three anonymous referees provided very helpful comments. Financial support by the Economic and Social Research Council (grant number ES/T015357/1) is gratefully acknowledged. For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission. The experimental protocol was approved by the Ethics Committee of the University of Essex, ETH2122-0287.

1 Introduction

Robo-advisors are financial tools that utilize algorithms to offer financial advice and automated trading tailored to investors' needs and preferences. As they only require minimal human intervention, they are inexpensive in comparison to traditional financial advice given by human experts. For this reason, they have also been praised for making financial advice available to investors for whom human advice is prohibitively expensive, thus having the potential to increase financial market participation (see e.g. Lieber, 2014). Despite the market for robo-advice being in its infancy, providers already have close to two trillion assets under management in 2025 (see e.g. Statista, 2025). Moreover, the industry has consistently featured robust growth rates and is expected to continue this growth path.

Robo-advisors differ from existing investment platforms and online brokers with respect to both customer assessment and customer portfolio management (Jung et al., 2018). Recent robo-advisors build on data gathered from investors to perform so called "risk profiling", such as surveys to elicit investment objectives, financial situation, and attitude towards risk (Bhatia et al., 2020). They then suggest or implement investment strategies that are the same for individuals who fall in the same category (D'Acunto and Rossi, 2019). Depending on the specific robo-advisor either a portfolio is suggested or is automatically implemented. Robo-advisors then typically rebalance the portfolio to keep the asset mix close to the optimum (D'Hondt et al., 2020; D'Acunto and Rossi, 2019). The precise internal workings of robo-advisors are usually proprietary and stay fairly opaque. For example, this is pretty much all Vanguard, a market leader in robo-advice, tells their customers about their robo: "When you sign up for a robo-advisor, it'll ask for basic information about your goals, risk tolerance, and the length of time you want to stay invested. Then technology takes over to suggest a portfolio for you . . . and it'll manage those investments over time, rebalancing periodically to make sure your asset mix stays on the right track. All behind the scenes, all automatically. It's pretty cool when you think about it." (Vanguard, 2023)

Given the increased usage of robo-advisors, we are interested in whether – and if so how – they can help investors make better decisions and whether they consequently have the potential to increase financial market participation. In order to study these questions we have conducted a large-scale, 10-week long, field experiment-like, pre-registered (AEA RCT Registry, AEARCTR-0009159) experiment.[1] The experimental approach is particularly

[1] https://www.socialscienceregistry.org/trials/9159

suited to our research question as it allows us to control for several key properties that would be difficult to control for in field data. Crucially, it allows i) to exogenously vary whether subjects have access to robo-advice or not and change key properties of robo-advice, ii) to measure individual risk attitudes of financial market participants (but also of non-participants) and based on these iii) to derive benchmarks for optimal behaviour. Most of this would be difficult if not impossible with field data.

Our experiment features two parts. In the first part, we elicit the risk preferences and some demographics from all subjects. Furthermore, we provide information about the second part, a financial market that subjects can choose to participate in. Depending on the treatment, subjects are also told that i) in that financial market they would receive advice by an algorithm fine-tuned to their risk preferences (treatment SOFT-ROBO), ii) an algorithm would implement their optimal investment decisions but they would retain the opportunity to overrule its choice (treatment HARD-ROBO), or iii) they would make decisions without the aid of an algorithm, and are not informed about its existence or usage by other participants (treatment CONTROL). Before we elicit risk preferences, we inform subjects in the robo treatments that their answers would be used to calibrate the algorithmic advice they receive.

After completing the first part, subjects are given the option to participate in a follow-up study that lasted for ten weeks.[2] We view the propensity to participate in this follow-up experiment (depending on assigned treatment) as an indicator for the potential of robo-advisors to increase financial market participation. In this context, we find that the presence of robo-trading does not influence subjects' decision to participate in the follow-up experiment. In each of our treatments almost 80% of the invited participants joined our follow-up study. In a certain way this is surprising as subjects retain the same level of control across treatments and should have appreciated the additional help which they could have chosen to ignore (in SOFT-ROBO) or override (in HARD-ROBO). On the other hand, subjects are (yet) not aware about how useful this additional help would turn out to be. As we elaborate below this will have a significant impact on *continued* market participation.

In the second part of the experiment, investors receive an initial endowment of experi-

[2] We deliberately chose a dynamic portfolio choice environment rather than a sequence of independent, static investment decisions because a central feature of robo-advice in practice is its emphasis on portfolio maintenance and rebalancing over time. A static setup, in which each period is independent, would largely abstract from this dimension and would therefore be less suited to studying the role of default implementation and rebalancing, which are key mechanisms in our experiment.

mental wealth and then have to allocate this wealth in each of the ten weeks across a set of hypothetical assets. This is a challenging task. Finding the optimal portfolio weights is far from trivial. They have to rebalance their portfolio from week to week. They have to realize certain hedging opportunities and be careful not to invest in dominated assets. In all of these tasks, investors can potentially benefit from following the advice of a robo.

Our experiment allows us to measure the usefulness of robo-trading compared to investing without assistance across several dimensions. In this context, we find that robo-advisors lead to i) significantly more optimal rebalancing of assets, ii) significantly more optimal hedging, and iii) a significant reduction in funds invested in the dominated assets. In all but the last dimension hard robos significantly outperform soft robos, which is in line with the well-known power of default settings or decision inertia (see e.g. Choi et al., 2004). In fact, our research design featuring an "intermediate" soft robo allows us to break apart the general usefulness of robos (HARD-ROBO vs CONTROL) that is purely due to advice (SOFT-ROBO vs CONTROL) and the part that is due to the avoidance of decision inertia (SOFT-ROBO vs HARD-ROBO). Our analysis reveals that both factors play an important role.[3]

The fact that investors benefit from using a robo is not really surprising.[4] What is surprising is *how much* they benefit in our experiment. We develop a methodology to quantify (in terms of differences in certainty equivalents) how much investors benefit from the robo and find that in treatment HARD-ROBO investors gain by almost 7.5 GBP relative to the control treatment, which seems a lot given the initial endowment of 50 GBP. These benefits should easily cover the rather low fees for robo-advisors (usually below 1%). We emphasize that our experiment allows us to provide these measurements in terms of expected utility, thus providing a more fine-tuned assessment of the welfare effects of robo-trading than the literature using field data, which typically focuses on expected or realized returns.

Given the unequivocally positive effects of robo-trading in our financial market, we finally investigate the implications for continued market participation. While robos do not influence initial participation as measured by the propensity of subjects to take part in our follow up experiment, we find clear evidence that they affect participation over time. In particular, we observe substantial differences in the extent to which investors remain

[3] Note that since the literature using field data only studies data from one specific given robo employed by the respective investment firm (mostly hard-robos), this comparison is another benefit of our experimental approach.

[4] After all, if investors follow the robo's advice, their decisions are optimal by construction.

invested rather than moving fully into the safe, cash-like asset. In the final round, the share of investors holding all of their wealth in cash is surprisingly high in the CONTROL treatment (around 38%), lower in SOFT-ROBO (around 25%), and close to zero in HARD-ROBO (0.4%).

These differences are mirrored by measures of platform interaction (i.e., interface visits), which capture both monitoring behaviour and active portfolio adjustment. Investors in HARD-ROBO access the interface more frequently than investors in SOFT-ROBO, who in turn do so more often than investors without robo-advice.[5] Taken together, these patterns indicate that robo-advisors reduce disengagement from risky asset markets over time, particularly when advice is implemented by default. Note that higher retention in risky asset markets may indirectly translate into higher overall market participation if fewer investors drop out over time.

Overall, the stark difference between soft and hard robo-advice suggests that the key benefit of robo-advisors may lie not only in the quality of advice, but in how advice is implemented. By shifting the default from inaction toward an optimized invested position, hard robo-advisors may mitigate behavioural frictions such as inertia or status quo bias, which are known to contribute to low market participation and high cash holdings.

The rest of the paper is structured as follows. In the following subsection we address the literature which is related to our study. In Section 3 we describe our experimental design. Section 4 provides details on the theoretical solution of our portfolio optimisation problem. Section 5 shows the results, and section 6 concludes.

2 Related Literature

Since robo-advice and robo trading are fairly recent phenomena the literature is still in its infancy.[6] One prominent topic in the (experimental) literature concerns algorithm aversion versus algorithm appreciation, i.e. whether people are more or less likely to follow the advice of an algorithm. Most of this literature focuses on a comparison of the reactions to advice given by humans and advice by an algorithm. Here the picture that arises is mixed. Some authors find support for algorithm aversion (e.g. Dietvorst et al., 2015,

[5] This is noteworthy because investors in HARD-ROBO could fully delegate portfolio decisions to the robo-advisor and thus engage less with the platform. A substantial fraction of visits involve confirming the hard robo's allocation rather than overriding it, see Section 5. Thus, engagement in HARD-ROBO should be interpreted primarily as monitoring behaviour rather than active portfolio re-optimisation.

[6] See D'Acunto and Rossi (2019, 2023) for surveys of the literature.

2018) while others provide indications that individuals appreciate algorithms (e.g. Logg et al., 2019). Also within the domain of financial decision making the evidence is mixed. In Holzmeister et al. (2022) investors are more inclined to delegate their decisions to an algorithm than to a human while in Germann and Merkle (2022) there is no strong evidence in either direction. Gaudeul and Giannetti (2025) find that algorithm adoption depends mainly on how successful traders were trading on their own in comparison to the case where algorithmic advice was available. Employing field data from a "hybrid"-robo-advising firm, where portfolio management is automated but investors are randomly matched with human advisors of varying quality, Greig et al. (2022) study ways in which algorithm aversion can be assuaged by human interaction. Our approach differs from the above by asking whether subjects are more inclined to participate in a financial market when an algorithm is available to the case when no advice (neither algorithmic nor human) is available. It can be argued that for financial decision making not having any advice at all is the more relevant benchmark as human advice is often orders of magnitude more expensive and thus not available to the kind of investors robo-trading appeals to. Concerning the question of algorithm aversion in this context, our paper somehow sits in the neutral camp; while we find no evidence that algorithms increase financial market participation there is some evidence that they fostered continued participation in our experiment.[7]

Further, our contribution differs from previous experimental ones by casting algorithmic advice in a financial setting where it is currently most prominently applied, namely optimal portfolio choice where investors have to first choose a portfolio that matches their risk profile and then rebalance this portfolio as profits and losses are realized so to keep it aligned to their risk preferences (see D'Acunto and Rossi 2019 for a more detailed description of applications and the internal workings of robo-advice). In contrast, Holzmeister et al. (2022) feature a setting with a series of independent lottery choices and no need for portfolio rebalancing. In Germann and Merkle (2022) and Gaudeul and Giannetti (2025) an unknown underlying state of the world determines payoffs and algorithms use Bayes rule to learn this state and may prevent individuals from overtrading. Moreover, whilst in Holzmeister et al. (2022) investors were asked to indicate their desired level of risk on a fixed scale, in our contribution investors' risk preferences were elicited. In contrast, in Germann and Merkle (2022) and Gaudeul and Giannetti (2025) the algorithms are not

[7] This latter aspect echoes findings of D'Acunto et al. (2019) in the field where adopters of robo-advice increase attention based on online account logins.

responsive to their client's risk attitude and instead simply maximise expected returns.[8]

The empirical literature has mainly been exploiting existing data sets from wealth management firms to identify the characteristics of investors who adopt robo-advice and to understand how portfolios evolve post adoption. Our findings about the positive effects of robo-advice are echoed in this literature. D'Acunto et al. (2019) show that investors exhibit less behavioural biases, such as the disposition effect, trend chasing, and the rank effect. The results regarding the disposition effect are confirmed in an experiment by Back et al. (2023) who compare a baseline without advice to several robo-advisor designs and show that robo-advisors can improve investment outcomes by mitigating behavioural biases such as the disposition effect. While their focus is on trading behaviour and bias reduction, our experiment complements this work by studying a dynamic portfolio choice problem with rebalancing and by distinguishing between non-binding advice and default implementation. This allows us to study a different channel, namely the role of defaults and inertia in sustaining investment over time.

Similarly, Rossi and Utkus (2024) find that robo-trading reduces idiosyncratic risk by lowering holdings of individual stocks and active mutual funds and raises exposure to low-cost indexed mutual funds. Moreover, Loos et al. (2020) document increased risk-taking (which may or may not be a good thing in terms of expected utility), holding more diversified portfolios with a larger fraction of index funds, and lower home bias and trend chasing tendencies, and an increase in (buy) turnover (which given fees is usually not a good thing). Our experimental approach allows us to add to this literature by characterizing welfare gains in terms of expected utility, something which would be fairly difficult to do with field data. Moreover, the empirical literature analyses behaviour in the face of existing commercial robo-advisors, the precise workings of which are often proprietary and not known to the researcher. In contrast, the experimental approach allows us to study how individuals interact with robos with known properties.

Furthermore, many of the empirical contributions only present correlational evidence as the investors decision to adopt robo-advice or not leads to observation of a selected sample. A notable exception is Bianchi and Briere (2026), who use a novel dataset from a large French employee savings plan. They find that employees who use a robo advisor

[8] In fact, not all financial institutions design their robos to account for their customers' risk preferences. For example, Scherer and Lehner (2021) show that the robo advisor they study makes recommendations mainly based on investment goals and time horizon. Thus, it is important to stress that our experiment is not necessarily applicable to all existing robo advisors but rather is about how an ideal robo advisor would work.

rebalance more and achieve higher returns. Given that the uptake of the robo advisor was voluntary, there is of course a potential endogeneity issue, which they address with unique feature of their data set, namely that they observe employees who show some interest in the robo but in the end do not sign up ("robo-curious"). Our experimental approach is a complementary solution to the endogeneity problem that allows to establish causation rather than correlation. It is noteworthy that those two very different approaches come to similar conclusions about the usefulness of robo advisors.

Another recent empirical paper by Reher and Sokolinski (2024) also finds a clever way of addressing the endogeneity issue of field studies. They use the fact that a bank lowered its account minimum balance for robo advice and show that consequently participation by middle-class investors increased. Compared to a control group, robo advice improved expected returns and yielded more diversified and personalized portfolios.

In a broader sense our paper also contributes to the (experimental) literature on delegation in financial decision making at large. Holzmeister et al. (2022) find that trust and blame shifting tendencies play an important role in the decision to delegate to human and robo-advisors. Moreover, with human advisors there may be a mismatch between the desired level of risk and the one implemented by the advisor. The importance of trust is echoed by Loos et al. (2019) who provide experimental evidence for Gennaioli et al.'s (2015) assertion that the higher the level of trust in a given money manager is, the more risk clients ask this advisor to take.

Another example where investors may delegate their financial decisions to others are copy trading platforms where investors on online platforms can copy each others' financial strategies. Similarly to robo-advising the cost of delegation is fairly low as compared to traditional advice. In contrast to robo trading or traditional financial advice, the delegatees are peers rather than algorithms or professional money managers. Apesteguia et al. (2020) show that in an experimental setting copy trading market a significant fraction of investors chooses to copy investors who have previously been successful by virtue of having taken on a lot of risk.[9] This chasing-past-performance tendency may lead to excessive risk taking in the aggregate. We speculate that similar forces could be at play in robo trading if investors can choose among several robos with different previous earnings.

[9] See also Freer et al. (2023) who study reasons why investors choose to copy/delegate in the first place.

3 Experimental Design

Our primary interest is in examining whether the use of robo-advice, either through recommendations or non-binding implementation, affects individuals' participation in financial markets, and how it influences their behaviour in terms of risk-taking and optimal asset selection. To this end, we conducted a large-scale, pre-registered (AEA RCT Registry, AEARCTR-0009159)[10] online experiment, which consisted of two parts.

For the first part, we invited a gender-balanced sample of 1,000 subjects on the platform Prolific who closely reflect the demographic characteristics of the general UK population with an average age of 39.7 and income levels roughly in line with UK median income (see Table B1 in the Appendix).[11]

During this part of the experiment, we provided subjects with an introduction to our financial market environment. Depending on their assigned treatment, subjects were told that i) they would either receive advice by an algorithm fine-tuned to their risk preferences (treatment SOFT-ROBO), ii) an algorithm would implement their optimal investment decisions but they would retain the opportunity to overrule its choice (treatment HARD-ROBO) or, iii) they would make decisions without the aid of an algorithm, and were not informed about its existence or usage by other subjects (treatment CONTROL). The number of subjects in the first part were 338 in CONTROL, 332 in SOFT-ROBO, and 330 in HARD-ROBO.

In the next step, we elicited subjects' risk preferences, which were used as the basis for providing algorithmic advice in the robo treatments. After completing the first part, subjects were given the option to participate in a follow-up study that lasted for ten weeks. In this second part of the experiment, they were asked to allocate their experimental wealth across multiple assets with varying expected payouts and levels of risk, including assets that were dominated by others. The first part of the experiment was conducted in May 2022 and transitioned directly into the second part. We now explain both parts of the experiment in more detail.

In the first part we administered a question from the Berlin Numeracy Test (BNT, see Cokely et al. 2012) before providing a description of the experimental market for the second part of the study. To encourage careful reading, we announced an upcoming quiz about the market. The description informed investors that they would receive an initial

[10] https://www.socialscienceregistry.org/trials/9159

[11] Median annual pay of all employees in the United Kingdom in 2022 was about 27,700 GBP (Office of National Statistics, 2022) which falls in our income category 3.

wealth of 50 GBP and could make one investment decision per week for the duration of the experiment. Each week, investors were presented with a set of assets to invest in, with each asset paying out according to a randomly determined state of the world. These assets had three possible states of realization, each with equal probability of occurring. We also informed them that we invite a total of 1,000 subjects to participate, and that we will randomly draw 100 individuals out of those who participate to be paid their actual portfolio value at the end of the experiment. In addition, we familiarized investors with the experimental market by providing an example of a state contingent asset (which was different to the ones used in the experiment).

In treatment CONTROL, investors received no additional information. In treatment SOFT-ROBO, investors were told that they would each week receive advice on their optimal portfolio choice by an algorithm and that these recommendations would account for their attitudes towards risk taking. We informed investors that their risk preferences would be measured through an upcoming lottery task, which was designed to calibrate the algorithmic advice they would receive. The description of treatment HARD-ROBO was similar, but investors were told that the algorithm would invest on their behalf while they would have the opportunity to override the proposed allocation if they wish to do so.[12] All investors then went through a short quiz on the general setup of the second part of the experiment before proceeding to the lottery task.

We used a simple incentivized lottery choice task to elicit subjects' risk preferences (cf. Binswanger 1980; Eckel and Grossman 2008). Subjects were asked to choose between five lotteries displayed in Table 1. Each lottery had two outcomes which occurred with probability $\frac{1}{2}$. The lotteries differ from each other with respect to mean payoff and payoff variance. Subjects in the robo treatments had been informed that their decision in this stage was relevant for the follow up experiment through calibrating the algorithmic advice.[13] According to their choice, we assigned each subject a representative coefficient of relative risk aversion γ and used this parameter for calculating optimal portfolio decisions (see next section). Given the lottery choices and under the (obviously restrictive) assumption of a CRRA utility function, we can impute coefficients of relative risk aversion. We

[12] The exact description can be found in the instructions in the online Appendix.

[13] For example, in treatment HARD-ROBO, the instructions said: "You may delegate your decision on how to invest to an algorithm (a "robo advisor") that invests and re-adjusts your current holdings across the available assets automatically. To allow the algorithm to understand your individual preferences, it needs to be calibrated."

chose the (rounded) mid points of the intervals that would yield the corresponding choice.[14]

Table 1: Lotteries for the Risk Elicitation Task

Lottery	Reward A	Reward B	Imputed γ
1	760	760	3
2	640	900	2
3	560	1040	0.8
4	320	1440	0.2
5	40	1660	0

Note: Subjects chose one of the 5 lotteries. The chosen lottery was played out for each subject and rewards A and B were chosen with probability 1/2 each. The exchange rate for rewards was 800 points to 1 GBP.

After the risk elicitation task, subjects proceeded to a questionnaire about financial market experience and familiarity with robo-trading. Finally, they were asked whether they would like to participate in our follow-up experiment. If they answered in the affirmative, they could register for the follow-up experiment by using a personal link. For those who did not, the experiment ended. Most subjects spent about 10 minutes on this part of the experiment. Every subject was paid and the median payoff was 1.58 GBP.

In the follow-up experiment, which was conducted over a 10-week time span from May to July 2022, subjects first had to register an e-mail address and provide informed consent. Subsequently they received a more detailed description of the financial market environment and, conditional on being in one of the robo treatments, the algorithmic advice. After these instructions, subjects had to answer four questions regarding key aspects of the financial market environment and were only allowed to proceed to the main stage if they had answered all of these questions correctly.[15] In this part of the experiment, 256 subjects participated in CONTROL, 242 in SOFT-ROBO, and 244 in HARD-ROBO.

We reiterated that it is their choice how to allocate funds across a set of available assets which yield returns depending on the realization of one of three equiprobable states of the world. We also emphasized that the number of available assets may change over time. In total, our experiment featured four different assets, A, B, C and D. During the first three

[14] For the unbounded top and bottom interval we had to make somewhat arbitrary choices. Furthermore, we had unfortunately a typo in the questionnaire regarding lottery 1. Fortunately, the optimal portfolio shares are very inelastic with respect to γ for highly risk averse investors such that the deviation is minor in quantitative terms (see Section C.1 in the Appendix for details).

[15] The exact instructions and quiz questions again can be found in Appendix C.

weeks, only assets A and B were available. From week four onward, investors also had asset C at their disposal. In the final four weeks, they could allocate among the full set of assets. Table 2 provides the gross returns in percent for each asset.

Asset A is a safe asset with negative real interest rate. Thus, holding exclusively asset A corresponds to essentially withdrawing from our financial market. One reason why we chose an asset with negative return was to discourage inaction by investors. Another reason was that it is simply realistic that holding cash has a negative real return in times of inflation. The other assets, in contrast, have higher expected returns, albeit at the cost of volatility.

Table 2: Assets

Asset	State 1	State 2	State 3	Weeks available
A	90	90	90	1-10
B	280	10	10	1-10
C	10	280	10	4-10
D	120	120	36	7-10

Note: States are drawn i.i.d. each week with equal probability.

If investors did not submit an allocation in a given week, then in treatments CONTROL and SOFT-ROBO, their funds would remain invested as they were at the beginning of the week, while in treatment HARD-ROBO the algorithmic advice would automatically be implemented.[16] At the end of each week a state of the world was drawn and the amounts invested in each asset would evolve according to this realization. For instance, if an investor had invested 20 GBP in asset A and 30 GBP in asset B, and state 2 realized, they would start the next week with an allocation of $20 \cdot 0.9 = 18$ GBP invested in asset A and $30 \cdot 0.1 = 3$ GBP invested in asset B. If investors did not change their portfolio, this allocation would roll over to the next week. Thus, in treatments CONTROL and SOFT-ROBO, there was a need for portfolio rebalancing; a risk-averse individual who chose not to do so would end up with a sub-optimal portfolio in the next week.

In treatment SOFT-ROBO, investors received non-binding advice on their optimal portfolio corresponding to their previously revealed risk category. In treatment HARD-ROBO, this advice was implemented as the investors' default portfolio in each week, but they could deviate from this implementation by overriding the robo's allocation. In all treatments, investors were sent weekly reminders to check their portfolio and make their

[16] In the first week, the entire endowment of subjects in CONTROL and SOFT-ROBO was invested in asset A.

choices.

The advice of the robos were based on optimal choices a CRRA investor would make given the elicited γ. However, investors were only told the following: "Your default investment is set by an algorithm that is supposed to support you when making your decisions. It was calibrated in the first stage of the experiment to account for your personal preferences." The exact derivation of the theoretical optimal solution can be found in the next section. Independent of individual risk preferences, our portfolio problem has the following properties: i) in weeks 4-10, investors who hold a positive fraction of asset A have to choose assets B and C in equal proportions[17] and ii) in weeks 7-10 no investor should invest in asset D as it is state-wise dominated, for example by a combination of assets A, B and C such as $(\frac{4}{10}, \frac{3}{10}, \frac{3}{10}, 0)$.[18]

We introduced the assets sequentially for several reasons. First, by starting with just two assets and adding more over time, we gradually increased the complexity of the investment decision. Second, this approach allowed us to cleanly separate our research questions regarding hedging and the dominated asset, namely whether in week 4 all risk averse investors would hedge by holding assets B and C in equal proportions and whether in week 7 investors would refrain from buying the dominated asset D.

After the tenth round, we conducted an additional questionnaire to elicit information on investors' financial market experience, education, profession, and income. Finally, we randomly selected 100 investors from all investors of the follow-up experiment, and paid them their portfolio in full through either Amazon vouchers or bank transfer. The average payoff of those selected was 20.38 GBP, which is slightly higher than the average final portfolio value of all investors (19.22 GBP).

4 Theoretical solution

In this section we derive the optimal way of investing in our experiment with 10 weeks and (up to) four assets. In general this is a rather complex problem as the optimal investment strategy may depend on current wealth and on the remaining investment horizon.[19] To

[17] The reason for this is that only risk-averse investors will chose asset A and any risk averse subject will combine assets B and C in equal proportions to hedge risk across states 1 and 2.

[18] Dominated assets in real financial markets could be e.g. managed funds which have higher fees than ETFs but are based on the same index.

[19] Capponi et al. (2022) study an even more complex problem when a robo advisor has to adapt to clients potentially changing risk preferences.

simplify, we shall make the assumption that investors have a CRRA utility function $u(w) = \frac{w^{1-\gamma}}{1-\gamma}$ with coefficient of relative risk aversion $\gamma \neq 1$. For CRRA utility functions it is well known that the solution to the intertemporal portfolio problem is given by the (myopic) optimisation in each week and is independent of wealth and of the investment horizon (Samuelson, 1969; Back, 2010). This logic applies regardless of whether additional assets may become available and also does not depend on the remaining time horizon.

Note first that our four-asset problem can, without loss of generality, always be reduced to the choice between one risky and one safe asset. This is obvious for weeks 1-3 when only assets A and B are available. Here we only have to distinguish between state 1 (with probability 1/3) and its complement. Once asset C becomes available in week 4, all risk averse investors will hold assets B and C with equal shares to hedge against states 1 and 2. Thus we can maximise over asset A and the mix $(\frac{1}{2}B, \frac{1}{2}C)$. Here we can distinguish between state 3 (with probability 1/3) and its complement. Risk neutral investors would be indifferent between assets B and C and any mix thereof. Risk loving investors would be indifferent between assets B and C. Finally, even when asset D becomes available, expected utility maximisers would not hold any asset D since it is dominated state-by-state e.g. by portfolio $2/5 * A + 3/5 * (\frac{1}{2}B, \frac{1}{2}C)$.[20]

Thus, we consider an investor with wealth w who decides for one week on a fraction f to be invested in a risky asset R and a fraction $1 - f$ to be invested in a sure asset S. There are two states of the world, 1 and 2, which occur with probabilities p and $1 - p$, respectively. The risky asset pays returns of r_1 and r_2 in these two states. The sure asset pays a return of s regardless of the state of the world. The expected utility of an investor who invests a fraction f of her wealth in the risky asset is consequently given by

$$p\left(w\left(f(1+r_1)+(1-f)(1+s)\right)\right)^{1-\gamma}+(1-p)\left(w\left(f(1+r_2)+(1-f)(1+s)\right)\right)^{1-\gamma}.$$

The first order condition with respect to f is characterized by

$$\frac{p(r_1-s)}{(f(1+r_1)+(1-f)(1+s))^{\gamma}} = -\frac{(1-p)(r_2-s)}{(f(1+r_2)+(1-f)(1+s))^{\gamma}}.$$

[20] This result would also hold for most non-expected utility theories, like cumulative prospect theory (Tversky and Kahneman, 1992), rank-dependent utility (Quiggin, 1982), and regret theories like Bell (1982) or Loomes and Sugden (1982) as they all respect state-wise dominance. It may possibly be violated by the representativeness heuristic and the availability heuristic as in the Linda problem (Tversky and Kahneman, 1982).

Rearranging gives

$$\frac{f(1+r_1)+(1-f)(1+s)}{f(1+r_2)+(1-f)(1+s)} = \left(\frac{p(s-r_1)}{(1-p)(r_2-s)}\right)^{1/\gamma} =: \Phi,$$

where Φ is the CE of the ratio of expected excess returns. Note that $\Phi < 1$ if $pr_1+(1-p)r_2 > s$, that is if the risky asset has a higher expected value than the sure asset. Solving for the optimal share of the risky asset f and noting that this share is bounded above by 1 (no borrowing constraint) yields

$$f = \min\left[1, \frac{(1-\Phi)(1+s)}{s-\Phi s - r_1 + \Phi r_2}\right].$$

Using this formula we can now calculate for each risk category the optimal mix between assets A and B (in weeks 1-3) and between A and $(\frac{1}{2}B, \frac{1}{2}C)$ in weeks 4-10, respectively (see Table 3 for the optimal portfolio weights). This is also the mix suggested by the robo-advisors.

Table 3: Optimal Share of Asset A by γ

γ	3	2	0.8	0.2	0
Weeks 1-3	0.98	0.97	0.92	0.66	0
Weeks 4-10	0.93	0.89	0.76	0.19	0

Note: In weeks 1-3, the residual is invested in asset B. In weeks 4-10, the residual is split equally across assets B and C.

4.1 How to calculate deviations from optimality

To evaluate the usefulness of robo-advice it is crucial to properly calculate deviations from optimality with and without advice. Note first that expected value or even ex post realized payoffs are not a reasonable way to compare performance as only risk neutral investors would be interested in maximising expected payoffs. The premise of robo-advice is that it can help investors to invest according to their own risk preferences. Since we cannot compare expected utilities across investors, we shall calculate for each investor the certainty equivalent (CE) of their chosen portfolio and compare it to the CE of the optimal portfolio. The difference then yields the loss in GBP that investors suffer from investing suboptimally.

While this idea is straightforward, there is a more subtle point about how to aggregate mistakes over the 10 weeks of the experiment. Let CE_t^{actual} be the certainty equivalent, per GBP of wealth, of the actually chosen portfolio in week t. We shall argue that the CE of the entire dynamic portfolio choice starting in week 1, $\overline{CE}_1$, is, for CRRA utility, equal to the product of the round-per-round CEs

$$\overline{CE}_1 = \prod_{t=1}^{T} CE_t^{actual}.$$

For the optimal CE, CE_t^{opt}, this follows easily from dynamic programming techniques (see Back, 2010, p. 221). For arbitrary non-optimal investments we could not find the result in the literature yet, so we decided to provide a proof in Appendix A.

Accordingly, we shall measure the loss in total certainty equivalent by

$$L := \left(\prod_{t=1}^{10} CE_t^{opt} - \prod_{t=1}^{10} CE_t^{actual} \right) w_1, \tag{1}$$

where w_1 is the initial endowment in week 1. Thus, L measures the ex ante expected loss in GBP incurred by an investor who deviates from the optimal portfolio in some or all of the weeks.

5 Results

For the analysis of the results we follow our preregistration unless noted otherwise.[21] The first question outlined there concerns market participation, i.e. the question whether investors are more willing to participate in our investment experiment if they are informed about the existence of a robo-advisor. This question is important since robo-advisors have been praised for democratizing financial advice (given their low fees) and potentially increasing financial market participation. We measure the willingness to participate in two ways: (1) whether they say they want to participate and (2) whether they back up this claim by providing their email address. Figure 1 depicts the share of investors who indicated their willingness to participate in the investment experiment by providing their email address. Fisher's exact tests ($p > 0.626$) show that participation rates do not differ across treatments.[22] [23]

[21] See AEA RCT Registry, AEARCTR-0009159.

[22] All tests reported in this paper are two-sided. Corresponding regression results including observable controls are reported in the Appendix (Table B3).

[23] The corresponding figures for the first participation measure and for a third (unregistered) measure, which counts how many subjects actually reached the investment stage, are shown in Figures B1 and

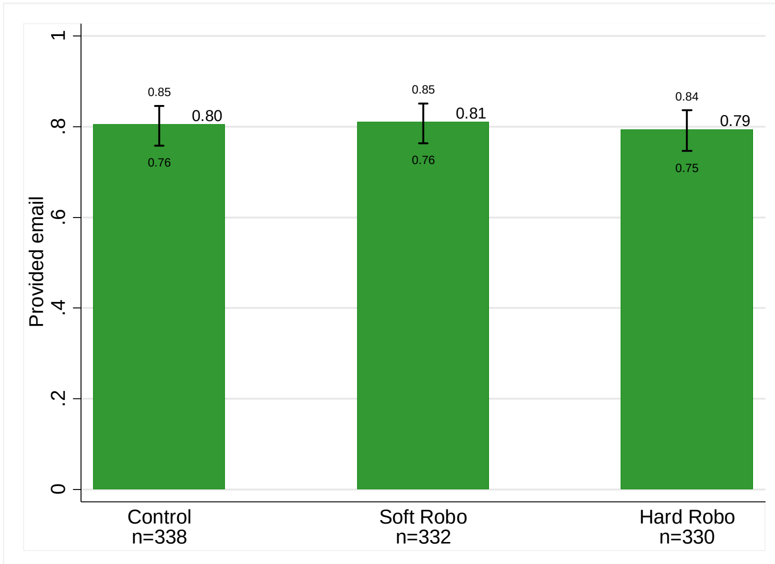

Figure 1: *Share of Subjects Providing Email*
Note: Shown are 95% Clopper–Pearson confidence intervals.

Result 1 *Offering a robo-advisor has no significant effect on initial financial market participation.*

Of course, we have to make sure that there is no selection of particular types into participation (and therefore treatments). Tables B1 and B2 in the Appendix show summary statistics for observables by treatment. All statistics seem to be very similar across treatments. In fact, we find no significant differences across treatments among those who participated in the investment experiment for any of the variables listed in Table B2 at the 10% level of t-tests.[24] For our most important variable, risk_category, there are no significant differences with p-values of $p > 0.310$.

While there does not seem to be a treatment effect on initial market participation, there might nonetheless be long-run effects on investors' decision to stay in the market. To study this question, we analyse the fraction of investors who invest *all* of their experimental wealth

B2 in the Appendix. None of the treatment differences are statistically significant, with the exception of a marginal difference in indicated willingness to participate between SOFT-ROBO and HARD-ROBO ($p = 0.086$). This isolated result is not reflected in the other participation measures.

[24]The only exception is the variable know_robo, likely because subjects in the robo treatments were informed about the existence of robo advisors right before the questionnaire.

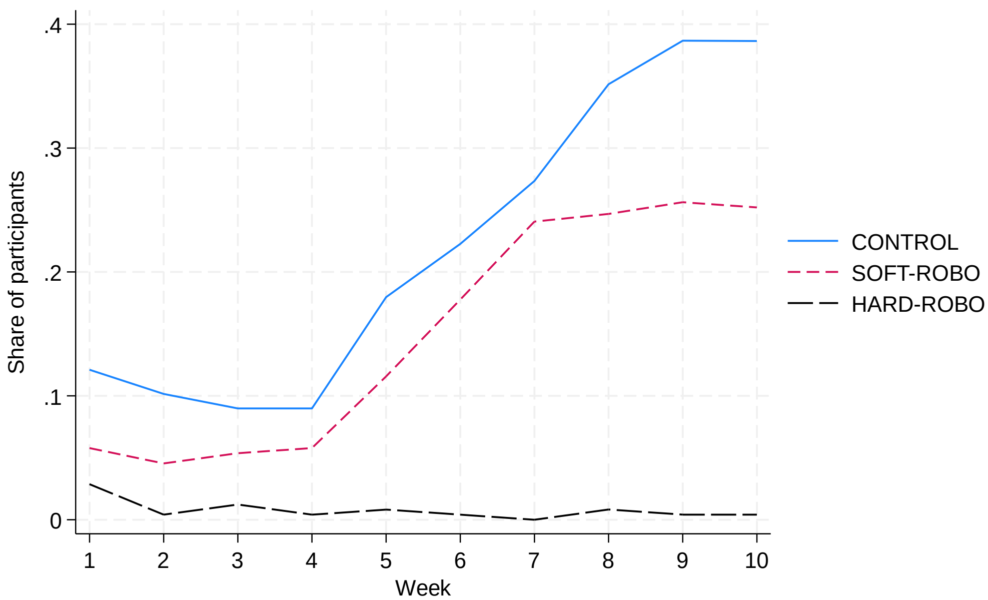

Figure 2: *Share of Investors in the Riskless Asset*

Note: For each treatment, we depict the share of investors who hold exclusively the riskless asset A.

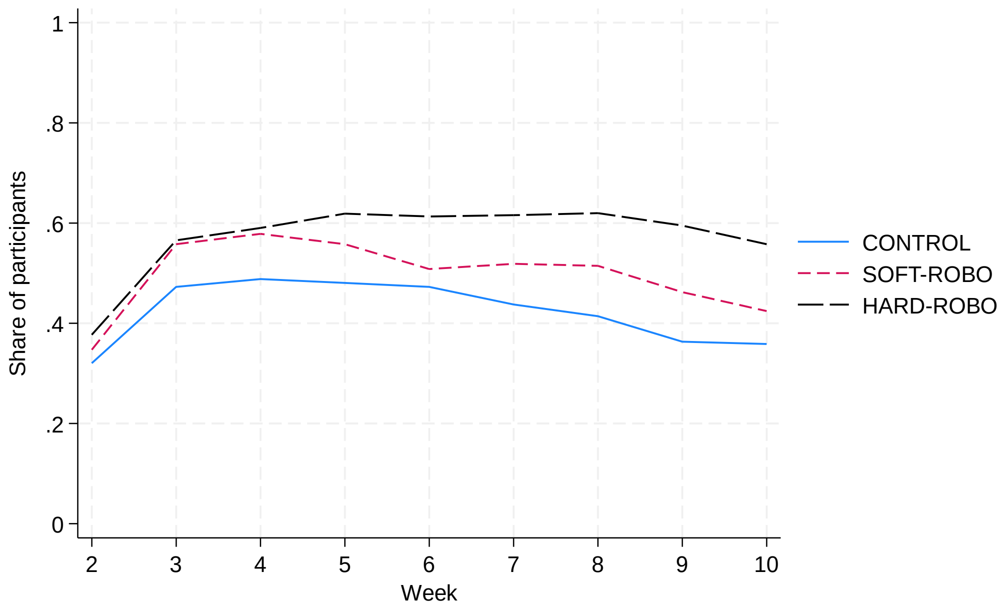

Figure 3: *Share of Investors on the Platform*

Note: For each treatment, we show the share of investors who visited and interacted with the investment platform. By default, everyone clicked in week 1. In week 2 there might have been a problem where some emails were caught in spam filters, which we fixed by week 3.

into asset A, the risk-less, negative return, cash-like asset. Since investors cannot withdraw their experimental wealth before the last round, this is as close as possible to staying out of financial markets altogether. Figure 2 shows the fraction of investors exclusively invested in asset A over all weeks, dropping those investors who were bankrupt.[25] The differences across treatments increase over time. In week 10, 38.65% of investors exclusively invested in asset A in CONTROL. In contrast, only 25.21% choose to do so in SOFT-ROBO, and a mere 0.41% do so in HARD-ROBO. All pairwise comparisons are statistically significant (Fisher exact tests: $p < 0.002$).

An alternative measure for continued financial market participation is provided by investors who actively engage with our financial market. We capture engagement by measuring the number of investors who clicked on the interface of our experiment in the respective week, again dropping those investors who were bankrupt. Note that this measure also captures monitoring behaviour and does not necessarily imply active portfolio adjustment. Figure 3 reports the results. Again, the treatment differences slightly increase over time. In the last week, 55.79% clicked in treatment HARD-ROBO, while 42.44% did so in SOFT-ROBO, and only 35.86% did so in CONTROL. Clicks are significantly more frequent in HARD-ROBO versus the other two treatments ($p < 0.004$), while the difference between CONTROL and SOFT-ROBO is not significant ($p = 0.139$).[26]

Result 2 *Robo-advisors significantly increase continued engagement with the investment environment, with Hard robos significantly outperforming Soft robos.*

Next, we come to our main question: Does a robo help investors? We shall split this question into several aspects that can be measured within our experiment. Does a robo prevent investing in a dominated asset? Does a robo encourage rebalancing of the portfolio? Does a robo help to hedge correctly? And finally, does the robo induce portfolios that are closer to the optimal ones given the preferences of investors?

Since asset D is dominated by a mix of the other assets state-by-state, no investor should ever choose to buy asset D. However, this fact is not obvious at all to many people,

[25] In the experiment 11 investors went bankrupt due to rounding of the balance. Only 2 of the 11 bankrupt investors come from HARD-ROBO.

[26] The high engagement rates in HARD-ROBO may be surprising at first sight as it by design does not require investors to engage with it. However it can be mainly explained by investors checking on their portfolio and/or confirming the robo's decision (roughly 66%). The remainder of investors (34%) engage with the platform in order to deviate from the robo's plan. Our high engagement rate is in line with D'Acunto et al. (2019) who, using field data from an investment platform, report increased engagement of investors who use a robo.

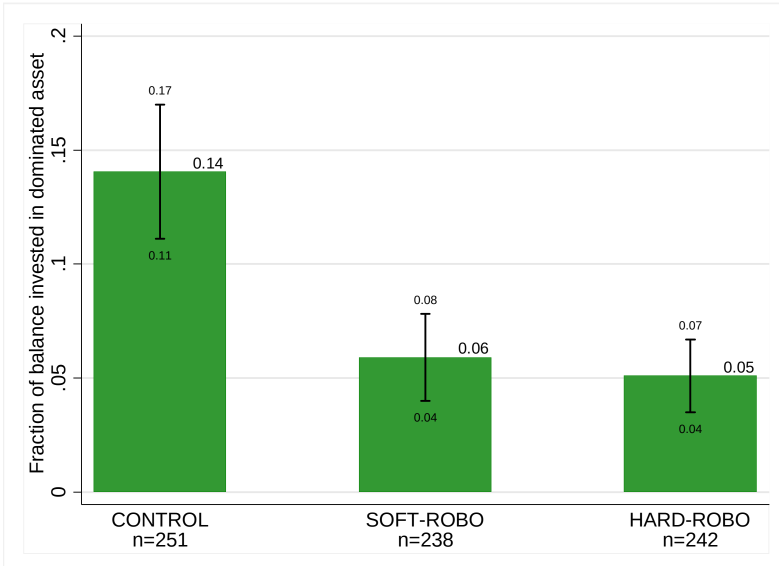

Figure 4: *Investment in the Dominated Asset*
Note: We depict fractions of balances invested in the dominated asset D by treatment. Shown are 95% confidence intervals.

in particular since asset D offers a higher payoff in two out of three states in any pairwise comparison to another asset. Figure 4 shows that investors in the CONTROL treatment invest about 14% of their balance into asset D in weeks 7-10. This is less than a $1/N$-rule would suggest but still substantial. In contrast, with the help of a robo-advisor, only slightly more than 5% of balances are invested in the dominated asset, regardless whether it is the Soft or the Hard robo. The difference is significant to CONTROL in both cases (t-test, $p < 0.001$). The reason why even in HARD-ROBO some funds are invested in the dominated asset is that some investors choose to override the robo.[27]

Result 3 *Both robo-advisors significantly reduce investment in dominated assets.*

One of the most important tasks of actual robo-advisors, which is prominently advertised by many service providers, is rebalancing. After the returns of the risky assets are drawn at the end of each week, investors have a new default portfolio, which will be

[27] One subject even sent us an email saying "Asset D is clearly the most attractive of the choices, so I guess it means that I'll be overruling the bot from now on unless it changes its tune."

carried over into the next week in treatments SOFT-ROBO and CONTROL unless they do something. In most cases, this default portfolio is not optimal given that the weights of the different assets have been shifted. We count how often investors rebalance in the right direction, which is defined as moving to a new portfolio which yields a higher certainty equivalent (CE_t^{actual}) than the default portfolio carried over from the previous week ($CE_t^{default}$). Figure 5 reports the share of portfolios that are rebalanced in the right direction. Since the Hard robo rebalances automatically, it is not surprising that this share is almost 1. In about 10% of weeks, however, investors actively override the robo and move to a portfolio with a lower certainty equivalent. In SOFT-ROBO, rebalancing in the right direction happens in only 40% of cases and in CONTROL in only 19% of cases. All pairwise differences are significant (using each investor as one independent observation, t-test, $p < 0.0001$).[28] The difference between HARD-ROBO and SOFT-ROBO is interesting and suggests that one reason why Hard robos are more successful is that investors often do not bother to adjust their portfolios when the gains of doing so seem small, although, of course, they add up. A major advantage of Hard robos is that they do not even require investors to engage with the platform every week.

Result 4 *With robo-advisors there is significantly more rebalancing, with Hard robos significantly outperforming Soft robos.*

Standard portfolio theory implies that, for risk-averse investors, risk should be diversified across assets whenever such diversification reduces portfolio risk without lowering expected returns. In our framework there is a particularly obvious opportunity to hedge away uncertainty, namely by holding assets B and C in equal proportions. In real life, hedging opportunities are usually much less obvious but this makes it even more surprising when investors without a robo often overlook this hedging opportunity in our experiment. Figure 6 shows the shares of successful hedging for the different treatments. Hedging is only required for risk averse investors. So we define a successful hedge if the weights of assets B and C do not differ by more than 10% or if an investor reveals risk neutrality or risk loving by not investing anything in riskless asset A.[29] All pairwise differences are

[28] Figure B3 in Appendix B reports a robustness check based on an individual-level measure that aggregates deviations from optimal rebalancing across weeks 2–10. The treatment comparisons remain unchanged.

[29] An alternative definition would be to require hedging for all investors classified as risk averse according to the risk elicitation task. This yields almost identical results.

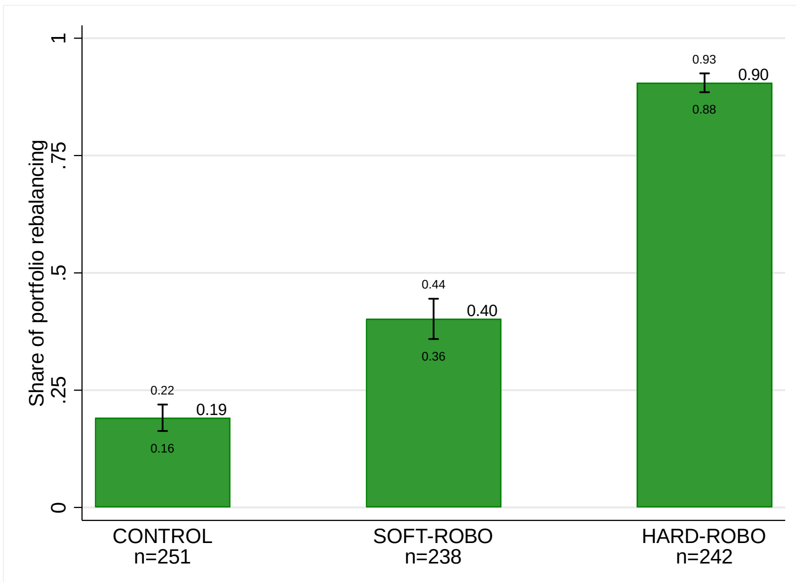

Figure 5: *Portfolio Rebalancing*

Note: We count how often investors rebalance in the right direction, which is defined as moving to a new portfolio that yields a higher certainty equivalent (CE_t^{actual}) than the CE of the default portfolio ($CE_t^{default}$) from the previous week. Calculated for weeks 2-10. We show shares by treatments and 95% confidence intervals.

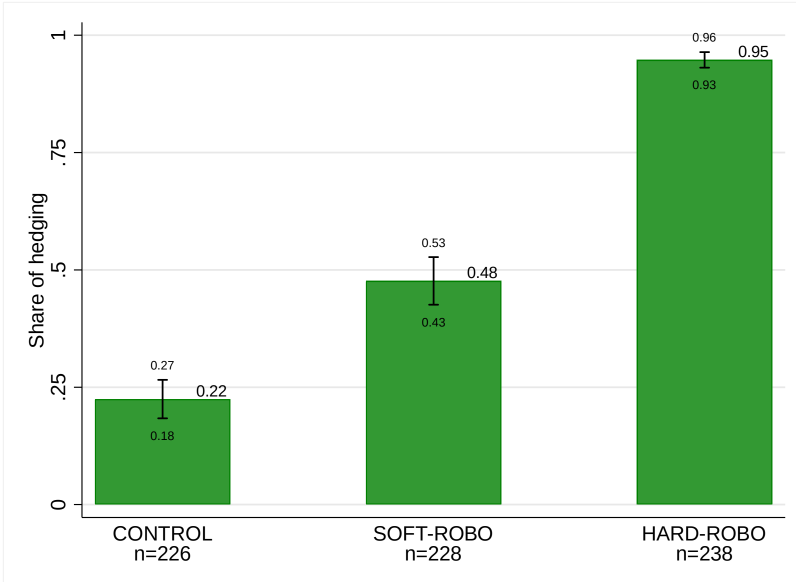

Figure 6: *Hedging*

Note: We depict shares by treatment. Correct hedges require that the weights of assets B and C do not differ by more than 10%. Excluded are 39 cases in which investors hold exclusively only or no asset A. Shown are 95% confidence intervals.

significant (using each investor as one independent observation, t-test, $p < 0.001$).[30]

Result 5 *With robo-advisors there is significantly more hedging, with Hard robos outperforming Soft robos.*

Finally, we can use the ex ante loss in certainty equivalent as defined in (1) as a monetary measure (in GBP) for how beneficial robos are for investors (see Figure 7). If investors had always followed the robos' advise, the loss would have been zero by definition. Due to overriding (in HARD-ROBO) or not following the advice (in SOFT-ROBO), investors have losses of 2.55 GBP and 5.37 GBP, respectively. But this is much less than the loss in CONTROL at 10.00 GBP. These losses are substantial as they amount to 20% of the initial endowment. All pairwise differences are significant (using each investor as one independent observation, t-test, $p < 0.001$).[31]

[30] The analysis with respect to hedging was not pre-registered and should be seen as exploratory.

[31] We excluded 11 bankrupt investors because we lack information on their hypothetical investment choices had they not been in bankrupcy. This makes it impossible to calculate $\overline{CE}_1$. See Figure B4 for a histogram

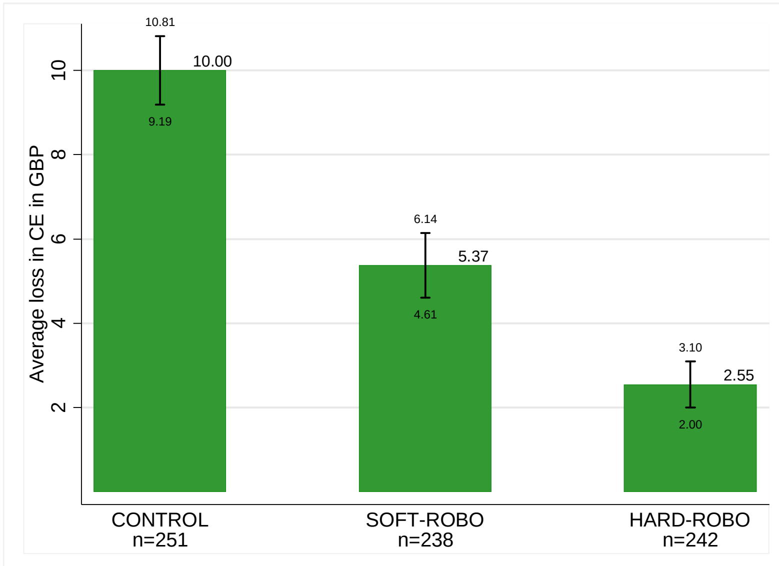

Figure 7: *Loss in Certainty Equivalent*

Note: By treatment, we show average loss in certainty equivalent measured in GBP together with 95% confidence intervals. We dropped 11 subjects who went bankrupt. Figure B4 shows the same analysis without dropping them.

Result 6 *The availability of a Robo avoids substantial losses in CE, with Hard robos again significantly outperforming Soft robos.*

This concludes the analysis specified in our preregistration. We now turn to an exploratory analysis aimed at better understanding how participants utilize robo-advice when it is available. While the treatment comparisons show substantial benefits of robo-advisors on portfolio outcomes, they also indicate that participants do not always adhere closely to the recommended allocations.

To study advice utilization in more detail, we introduce a continuous measure of advice following that captures how closely participants' portfolio choices align with the robo's recommendation in each week.[32] We compute for each participant and week the squared deviation between realized portfolio weights and the theoretically optimal portfolio weights. In our environment, the robo recommendation coincides with the theoretically optimal

including bankrupt investors.

[32] Our measure builds on the rebalancing measure proposed by Litterscheidt and Streich (2020).

portfolio, so lower values of this measure indicate closer adherence to the provided advice.

Formally, for each participant and week t, advice deviation is defined as

$$\Delta_t := \sum_{x=1}^{4} \left(f_x^t - \zeta_x^t\right)^2$$

where f_x^t denotes the share invested in asset x in week t, and ζ_x^t denotes the corresponding optimal share.

Using this measure, we examine differences in advice adherence across the SOFT- and HARD-ROBO treatments and explore how deviations from advice are associated with participants' current decision environment.

So why do investors refuse to follow the advice (in SOFT-ROBO) or even override it (in HARD-ROBO)? To shed some light on this question, we estimate linear regressions (see Table 4) with advice deviation Δ_t as the dependent variable. A natural behavioural motivation is that investors may experience regret after delegating decisions to the robo-advisor and subsequently realizing unfavourable outcomes. However, regret cannot be directly identified in our setting, as we do not observe the counterfactual portfolios investors would have chosen in the absence of robo-advice.

Rather than attempting to measure regret directly, we therefore focus on objectively unfavourable investment realizations. In particular, we use the occurrence of State 3, an outcome in which all risky assets generate losses, as a proxy for a clearly disappointing outcome. For these regressions, we restrict attention to investor-week observations in which participants held some amount of risky assets in the previous week, as otherwise no such disappointment can arise.

Table 4 shows that deviations from advice are smaller in HARD-ROBO compared to SOFT-ROBO, which is partly mechanical due to the default allocation in each treatment. However, the effect is still significant when conditioning on participants who actively interacted with the investment platform in that week. In the HARD-ROBO treatment, the occurrence of State 3 in the previous week is associated with larger subsequent deviations from the robo's recommendations, as is the presence of a dominated asset in later weeks. These effects are much smaller and not significant in treatment SOFT-ROBO.[33]

[33] Table B4 in Appendix B further extends this exploratory analysis by including advice deviation in the previous week and interaction terms. Table B5 provides complementary evidence using a binary definition of advice following (within a 10% band). Throughout, the occurrence of State 3 in the previous week has a significant effect on advice deviation in the HARD-ROBO treatment.

Table 4: Linear Regression: Why do Investors Deviate from the Robo's Advice?

	(1) Unconditional (Pooled)	(2) Engaged (Pooled)	(3) SOFT-ROBO (Engaged)	(4) HARD-ROBO (Engaged)
HARD-ROBO	-0.10^{***}	-0.05^{***}		
	(0.02)	(0.02)		
state 3 last week			0.01	0.05^{***}
			(0.02)	(0.02)
week with dominated asset			0.02	0.07^{**}
			(0.04)	(0.03)
week			0.02^{*}	0.01
			(0.01)	(0.01)
balance			0.00	0.00
			(0.00)	(0.00)
risk category			-0.01	-0.01
			(0.02)	(0.01)
numeracy			-0.03	0.03
			(0.03)	(0.02)
age			0.00	-0.00
			(0.00)	(0.00)
female			0.02	0.01
			(0.03)	(0.02)
constant	0.16^{***}	0.14^{***}	-0.01	-0.02
	(0.02)	(0.02)	(0.10)	(0.06)
Obs.	4,839	2,813	1,002	1,233

Note: The dependent variable is the squared distance between the participant's portfolio and the robo-advisor's recommended (optimal) portfolio in a given week, with lower values indicating closer adherence to advice. Columns (2)–(4) condition on weeks in which participants actively interacted with the platform. Column (1) is unconditional. Columns (3) and (4) use only observations in which investors held some amount of risky assets in the previous week. Standard errors are clustered at the participant level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6 Conclusion

Robo-advisors have been praised as important new tools for financial retail markets. In this paper, we studied their potential benefits in a large-scale portfolio choice experiment running over ten weeks. We find that investors' utility increases substantially when they have access to a robo-advisor. These benefits are significantly larger in the HARD-ROBO treatment, where robo-advice is implemented automatically, than in the SOFT-ROBO treatment, where investors merely receive recommendations and must implement them themselves. This is in line with existing evidence on the power of default settings (see e.g. Madrian and Shea, 2001; Choi et al., 2004).[34] Although the institutional setting and decision margins differ substantially, the common theme is that default options can have powerful effects on financial behaviour by reducing the need for active decisions. Our results suggest that similar forces may operate in investment environments beyond retirement savings, particularly when default implementation affects the propensity to remain invested over time. In contrast to existing field studies, our experimental design allows us to disentangle the benefits arising from improved advice that is present in both robo treatments from the benefits of overcoming investor inertia, which are specific to hard robo-advisors.

A striking result of our experiment concerns cash holdings. We observe a sharp reduction in investors holding exclusively cash in the HARD-ROBO treatment, in contrast to both the SOFT-ROBO and CONTROL treatments. This contrast is noteworthy given that European households are widely documented to hold substantial shares of their financial wealth in low-yielding cash and cash-like assets.[35] In our experimental setting, investors who receive no advice or only non-binding advice are substantially more likely to move fully into cash over time, whereas default implementation in the HARD-ROBO treatment largely prevents such behaviour. Since the SOFT- and HARD-ROBO treatments face identical risk and return profiles and the same degree of opacity regarding the algorithm, these differences cannot be explained by ambiguity aversion alone. Instead, default implementa-

[34] It is also in line with Bhattacharya et al. (2012) who show that advice overall has no significant effect for a group of customers who received portfolio advice from their bank. But it improved efficiency for those who follow the advice.

[35] For example, data for Europe indicate that roughly half of household wealth is held in bank accounts, with only about one-third of European households owning stocks (Financial Times, 2024). UK evidence from the FCA's Financial Lives Survey shows that among adults with investible assets of at least £10000, more than 40% hold all of their money in cash savings products rather than in investments (Financial Conduct Authority, 2024, p. 41).

tion appears to play a central role, consistent with evidence on inertia and status quo bias in financial decision-making. In this sense, hard robo-advisors appear particularly effective at mitigating behavioural frictions that lead investors to move fully into cash over time.

At the level of portfolio construction, we find that the benefits of robo-advisors materialize through several concrete features: (1) selecting and tailoring initial portfolios according to investors' risk preferences, (2) rebalancing portfolios over time, (3) avoiding dominated assets, and (4) using mixtures of assets to hedge risk across states. Many real-world robo-advisors are designed to do just that (Vanguard, 2023). Yet arguably, the relative importance of these features is likely to depend on the nature of the financial market under consideration.[36]

Existing empirical work on robo-advisors typically evaluates performance in terms of expected or realized returns. Robo-advisors, however, are designed to maximise expected utility, and only risk-neutral investors would be interested in maximising expected returns. We therefore compare certainty equivalents of portfolios with and without robo-advice. We find that hard robos help investors save almost 7.5 GBP on average out of an initial endowment of 50 GBP, which seems like a substantial gain, in particular given the generally low fees of robo-advisors.

Our experimental environment is intentionally stylized and abstracts from several features of real-world financial markets, including large asset universes, complex correlation structures, trading costs, fees, liquidity constraints, and rebalancing frictions. As a result, our findings should not be interpreted as providing quantitative predictions for portfolio performance or welfare gains in actual markets. Instead, the contribution of the experiment lies in isolating behavioural and institutional mechanisms through which robo-advice affects investment behaviour. In particular, the results highlight the importance of default execution in overcoming inertia, promoting disciplined rebalancing, and avoiding dominated investment choices. While these mechanisms are illustrated in a simplified setting, similar challenges arise, often in less transparent form, in real-world investment environments with many assets and greater complexity. Our results therefore speak to the direction and nature of these effects rather than to their precise magnitude.

Overall, robo-advisors appear to be an attractive alternative to traditional financial advice. For many small investors, the cost of traditional human advice is prohibitive

[36] For instance, the existence of dominated assets or hedging opportunities may be less transparent in some real markets while playing a more prominent role in others. Excessive fees can render financial instruments effectively dominated, whereas hedging opportunities may arise through internationally diversified ETFs.

anyway. But even when we abstract from cost, advice given by humans is often plagued by the same biases (see e.g. Barberis and Thaler 2003 or Hirshleifer 2015) as encountered with individual investors (Linnainmaa et al., 2021) and financial professionals when trading on behalf of clients often display no better decision quality than their clients (Stefan et al., 2022). Moreover, human advisors may be incentivized to recommend more trades than necessary to generate commissions (Hackethal et al., 2012). Whether robo-advisors can overcome such issues remains to be seen in future research.

On a more cautious note, as the prevalence of robo-advisors increases, it will be more important to understand their implications for financial markets at large. For instance, there is an ongoing discussion on the potential effects of algorithms on collusion in competitive settings (see e.g. Calvano et al. 2020 and Assad et al. 2024). Since the current generation of robo-advisors chooses portfolio allocations based on estimates from historical data, there currently does not seem to be much scope for collusion. However, future generations may rely on ever increasingly sophisticated models (such as BloombergGPT outlined in Wu et al. 2023). Thus, it may become more important to understand the potential of such algorithms to generate collusive outcomes and disrupt financial markets.

Author Affiliations

Marco Lambrecht: University of Stavanger
Joerg Oechssler: University of Heidelberg
Simon Weidenholzer: University of Essex

References

Apesteguia, J., Oechssler, J., and Weidenholzer, S. (2020). Copy trading. *Management Science*, 66(12):5608–5622.

Assad, S., Clark, R., Ershov, D., and Xu, L. (2024). Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. *Journal of Political Economy*, 132(3):723–771.

Back, C., Morana, S., and Spann, M. (2023). When do robo-advisors make us better investors? the impact of social design elements on investor behavior. *Journal of Behavioral and Experimental Economics*, 103:101984.

Back, K. (2010). *Asset pricing and portfolio choice theory*. Oxford University Press.

Barberis, N. and Thaler, R. (2003). A survey of behavioral finance. In Constantinides, G., Harris, M., and Stulz, R., editors, *Handbook of the Economics of Finance*, volume 1, pages 1051–1120. Elsevier Science.

Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, 30(5):961–981.

Bhatia, A., Chandani, A., and Chhateja, J. (2020). Robo advisory and its potential in addressing the behavioral biases of investors - A qualitative study in Indian context. *Journal of Behavioral and Experimental Finance*, 25:100281.

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. *The Review of Financial Studies*, 25(4):975–1032.

Bianchi, M. and Briere, M. (2026). Human-robot interactions in investment decisions. *Management Science*, 72(1):14–31.

Binswanger, H. P. (1980). Attitudes toward risk: Experimental measurement in rural India. *American Journal of Agricultural Economics*, 62(3):395–407.

Calvano, E., Calzolari, G., Denicolo, V., and Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10):3267–3297.

Capponi, A., Olafsson, S., and Zariphopoulou, T. (2022). Personalized robo-advising: Enhancing investment through client interaction. *Management Science*, 68(4):2485–2512.

Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2004). For better or for worse: Default effects and 401 (k) savings behavior. In Wise, D. A., editor, *Perspectives on the Economics of Aging*, pages 81–126. University of Chicago Press.

Cokely, E. T., Galesic, M., Schulz, E., Ghazal, S., and Garcia-Retamero, R. (2012). Measuring risk literacy: The Berlin numeracy test. *Judgment and Decision making*, 7(1):25–47.

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

D'Acunto, F. and Rossi, A. G. (2019). Robo-advising. In Rau, R., Wardrop, R., and Zingales, L., editors, *The Palgrave Handbook of Technological Finance*. Springer.

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.

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

Dietvorst, B. J., Simmons, J. P., and Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1):114–126.

Eckel, C. C. and Grossman, P. J. (2008). Men, women and risk aversion: Experimental evidence. In Plott, C. R. and Smith, V. L., editors, *Handbook of Experimental Economics Results*. Elsevier.

Financial Conduct Authority (2024). Financial lives survey 2024. https://www.fca.org.uk/publication/financial-lives/fls-2024-cash-savings.pdf (last accessed: 30 January 2026).

Financial Times (2024). Europe's savers keep too much money in cash. *Financial Times*. https://www.ft.com/content/860c1114-0812-4b82-8e3d-9a94fea2eb40 (last accessed: 30 January 2026).

Freer, M., Friedman, D., and Weidenholzer, S. (2023). Motives for delegating financial decisions. *Available at arXiv:2309.03419*.

Gaudeul, A. and Giannetti, C. (2025). Beyond performance: Exploring trade-offs in the design of financial algorithms. *Journal of Behavioral and Experimental Finance*, page 101119.

Gennaioli, N., Shleifer, A., and Vishny, R. (2015). Money doctors. *The Journal of Finance*, 70(1):91–114.

Germann, M. and Merkle, C. (2022). Algorithm aversion in delegated investing. *Journal of Business Economics*, pages 1–37.

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*.

Hackethal, A., Haliassos, M., and Jappelli, T. (2012). Financial advisors: A case of babysitters? *Journal of Banking & Finance*, 36(2):509–524.

Hirshleifer, D. (2015). Behavioral finance. *Annual Review of Financial Economics*, 7:133–159.

Holzmeister, F., Holmén, M., Kirchler, M., Stefan, M., and Wengström, E. (2022). Delegation decisions in finance. *Management Science*, pages 4363–4971.

Jung, D., Dorner, V., Glaser, F., and Morana, S. (2018). Robo-advisory: digitalization and automation of financial advisory. *Business & Information Systems Engineering*, 60:81–86.

Kahneman, D. and Tversky, A. (1982). On the study of statistical intuitions. *Cognition*, 11(2):123–141.

Lieber, R. (2014). Financial advice for people who aren't rich. *The New York Times*, April 12, 2014:B1.

Linnainmaa, J. T., Melzer, B. T., and Previtero, A. (2021). The misguided beliefs of financial advisors. *The Journal of Finance*, 76(2):587–621.

Litterscheidt, R. and Streich, D. J. (2020). Financial education and digital asset management: What's in the black box? *Journal of Behavioral and Experimental Economics*, 87:101573.

Logg, J. M., Minson, J. A., and Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151:90–103.

Loomes, G. and Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92(368):805–824.

Loos, B., Germann, M., and Weber, M. (2019). Trust and delegated investing: A money doctors experiment. *mimeo*.

Loos, B., Previtero, A., Scheurle, S., and Hackethal, A. (2020). Robo-advisers and investor behavior. *Unpublished Working Paper*.

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

Office of National Statistics (2022). Employee earnings in the UK: 2022. https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours /bulletins/annualsurveyofhoursandearnings/2022 (last accessed: 26 April 2026).

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization*, 3(4):323–343.

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

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

Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, 51(3):239–246.

Scherer, B. and Lehner, S. (2021). What drives robo-advice? *Available at SSRN 3807921*.

Statista (2025). Robo-advisors: Assets under management. https://www.statista.com/forecasts/1262614/robo-advisors-managing-assets-worldwide (last accessed: 12 December 2025).

Stefan, M., Holmén, M., Holzmeister, F., Kirchler, M., and Wengström, E. (2022). You can't always get what you want: An experiment on finance professionals' decisions for others. *Available at SSRN 4018654*.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323.

Vanguard (2023). What is a robo-advisor? https://investor.vanguard.com/advice/robo-advisor (last accessed: 5 July 2023).

Wu, S., Irsoy, O., Lu, S., Dabravolski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D., and Mann, G. (2023). BloombergGPT: A large language model for finance. *arXiv preprint arXiv:2303.17564*.