HEIDELBERG UNIVERSITY
DEPARTMENT OF ECONOMICS

UNIVERSITÄT HEIDELBERG
ZUKUNFT SEIT 1386

On the benefits of robo-advice in financial markets

Marco Lambrecht

Jörg Oechssler

Simon Weidenholzer

AWI DISCUSSION PAPER SERIES NO. 734
September 2023

On the benefits of robo-advice in financial markets*

Marco Lambrecht, Jörg Oechssler, and Simon Weidenholzer†

September 13, 2023

Abstract

Robo-advisors are novel tools in financial markets that provide investors with low–cost financial advice, usually based on individual characteristics like risk attitudes. In a portfolio choice experiment running over 10 weeks, we study how much investors benefit from robo advice. We also study whether robos increase financial market participation. The treatments are whether investors just receive advice, have a robo making all decisions for them, or have to trade on their own. We find no effect on initial market participation. But robos help investors to avoid mistakes, make rebalancing more frequent, and overall yield portfolios much closer to the utility maximizing ones. Robo-advisors that implement the recommendations by default do significantly better than those that just give advice.

JEL codes: C91, D81, G12, G20, G41.
Keywords: algorithmic trading, experiment, financial markets.

*We would like to thank seminar audiences at Ben Gurion University, Bern, Durham, Helsinki, Nottingham, WU Vienna, Heidelberg, ESA Bologna, SEET Valancia, and the Experimantal Finance Conference in Sofia for comments. We are also thankful to Sebastian Ebert, Jean Paul Rabanal, and Martin Weber for very useful discussions. Financial support by the Economic and Social Research Council (grant number ES/T015357/1) is gratefully acknowledged.

† `Lambrecht`: Hanken School of Economics & Helsinki GSE, Arkadiankatu 7, PO Box 479, Fi-00101 Helsinki, email: marco.lambrecht@hanken.fi, `Oechssler`: Department of Economics, University of Heidelberg, Bergheimer Str. 58, 69115 Heidelberg, Germany, email: oechssler@uni-hd.de. `Weidenholzer`: Department of Economics, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, United Kingdom, email: sweide@essex.ac.uk.

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 in excess of two trillion assets under management (see e.g. Statista, 2023). 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 re-balance the portfolio by buying and selling assets at fixed sampling intervals, mostly quarterly (D'Hondt et al., 2020; D'Acunto and Rossi, 2019). The actual trading strategies for the various categories are usually developed in research departments of the robo-advisors and rely on estimates of expected returns and co-variances obtained from historical data. The precise internal workings of robo-advisors are usually proprietary and stay fairly opaque as the following quote from one of the leading robo-advice providers *Vanguard* illustrates: "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, pre-registered (AEA RCT Registry, AEARCTR-0009159) online experiment. The experimental approach is particularly 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 behavior. Most of this would be difficult if not impossible with field data.

Our experiment features two parts. The first part, subjects receive some information about the financial market, in which they can participate in the second part. Depending on the treatment, subjects are also told that i) 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). After this we elicit risk preferences, informing subjects in the robo treatments that their answer 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. 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 (either through non-binding advice in SOFT-ROBO or hard wired in HARD-ROBO) 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, subjects initially receive an endowment of experimental wealth. They then have to allocate this wealth each week across a set of assets over the period of ten weeks. Among these assets there is always a riskless asset which depreciated in value every round, thus capturing the effect of holding cash in times of inflationary pressures. Initially, there is one risky asset with higher return and over time additional assets are added to the set of available assets. Our set of assets matched several properties of real financial markets -some of which may have not been obvious for investors on first sights- for which a robo-advisor might be really useful: i) In each round the starting portfolio would be given by the realized portfolio in the previous period and so would generally deviate from the optimal portfolio given the subjects' risk preferences, thus requiring re-balancing of portfolios. ii) In some rounds, buying a certain combination of assets would allow investors to hedge away some of the risk at no extra cost. iii) In some rounds, there was an asset which appeared attractive on first inspection but turned out to be dominated by a combination of other assets. In all of these dimensions, subjects can benefit from following the advice of a robo (in SOFT-ROBO) or not deviating from the optimal trading strategy implemented (in HARD-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 fact, in all but the last dimension hard robos significantly outperform soft robos, which is in line with the well-known power of default settings (see e.g. Choi et al., 2004).

The fact that investors benefit from using a robo is not really surprising.[1] What is surprising is how much they turn out to benefit in our experiment. We develop a methodology to quantify (in terms of differences of 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 seem to influence initial market participation as measured in the propensity of subjects to take part in our follow up experiment, there is clear evidence that they foster continued participation in the sense of the number of investors engaging with our financial market.

[1] Note that if subjects follow the robo's advice, their decisions are optimal by construction.

In particular, we find that the share of investors parking all of their money in the safe, cash like, asset in the final round is surprisingly high (around 38%) in the control treatment. It is significantly lower (around 25%) in SOFT-ROBO and close to zero in HARD-ROBO (0.4%). This finding is echoed when analyzing engagement rates, i.e. the number of investors clicking on the interface. Investors with the hard robo significantly more often engage with the platform than investors with the soft robo who in turn significantly more often engage than investors without a robo.[2] Note that the positive effects of continued participation in financial markets may indirectly lead to an increase in overall financial market participation as over time, fewer investors drop out.

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 2 we describe our experimental design. Section 3 provides details on the theoretical solution of our portfolio optimization problem. Section 4 shows the results, and section 5 concludes.

1.1 Related Literature

Since robo-advice and robo trading are fairly recent phenomena the literature is still in its infancy.[3] The empirical literature has mainly been exploiting existing data sets from wealth management firms to identify the characteristics of investors (e.g. demographics and current portfolio properties) who adopt robo-advice (D'Acunto et al., 2019) or hard-wired robo trading (Rossi and Utkus, 2020) and to understand how portfolios evolve post adoption.

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 impact of advice given by humans and the impact of advice by an algorithm. Here the picture that arises is mixed. Some authors find support for algorithm aversion (e.g. Dietvorst et al., 2015) 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

[2] This is even more surprising because, by design, subjects in treatment HARD-ROBO had the opportunity to delegate the decision to the robo advisor and thus could have adopted a lower engagement rate compared to the other treatments while still maintaining optimal portfolio decisions. It is worth to note that their higher propensity to click is not driven by the urge to overwrite the robo, as most of these investors do not deviate from the robo recommendation. For details, see Section 4.

[3] See D'Acunto and Rossi (2019) for a discussion of this early literature

algorithm than to a human while in Germann and Merkle (2022) there is no strong evidence in either direction. Gaudeul and Giannetti (2023) 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. 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.[4]

Further, our contribution differs from previous ones in 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 re-balance 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 (2023) 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 where 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 (2023) the algorithms are not responsive to their client's risk attitude and instead simply maximize expected returns.

Our findings about the positive effects of robo-advice are echoed in the empirical literature employing field data. D'Acunto et al. (2019) show that investors exhibit less behavioral biases, such as the disposition effect, trend chasing, and rank effect. Similarly, Rossi and Utkus (2020) 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

[4] 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.

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.

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.[5] 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.

2 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 behavior in terms of risk-taking and optimal asset selection. To this end, we conducted a large scale, pre-registered (AEA RCT Registry, AEARCTR-0009159) 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

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

with an average age of 39.7 and income levels roughly in line with UK median income (see Table 4 in the appendix).[6]

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). 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 pay-outs and levels of risk, including assets that were dominated by others. 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 subjects that they would receive an initial wealth of 50 GBP and could make one investment decision per week for the duration of the experiment. Each week, subjects 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 subjects 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, subjects received no additional information. In Treatment SOFT-ROBO, subjects 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 subjects that their risk preferences

[6]Median annual pay of employees in the United Kingdom in 2022 was about 27,000 GBP (Office of National Statistics, 2022) which falls in our income category 3.

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 subjects 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.[7] All subjects 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 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.[8] 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 chose the (rounded) mid points of the intervals that would yield the corresponding choice.[9]

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. Rewards A and B were chosen with probability 1/2 each.

After the risk elicitation task, subjects proceeded to a questionnaire about financial

[7] The exact description can be found in the instructions in the online appendix.

[8] For example, in Treatment HardRobo, 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."

[9] 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 5.1.4 in the appendix for details).

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.[10]

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 weeks, only assets A and B were available. From week four onward, subjects 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. The other assets, in contrast, have higher expected returns, albeit at the cost of volatility. Thus, holding exclusively asset A corresponds to not participating in our financial market in times of inflationary pressures.

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 period with equal probability.

Importantly, subjects were informed about the (treatment dependent) evolution of their investment if they do not submit an allocation in a given week. In CONTROL and SOFT-

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

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.[11] After each period a state of the world was drawn, and in treatments CONTROL and SOFT-ROBO the amount invested in each asset would evolve according to this realization.[12] If subjects did not change their portfolio, the allocation would roll over to the next period. Thus, in these treatments, 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 period. In treatment SOFT-ROBO, subjects 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 subjects' default portfolio in each week, but they could deviate from this implementation by overriding the robo's allocation. In all treatments, subjects were sent weekly reminders to check their portfolio and make their choices.

The advice of the robos were based on optimal choices a CRRA investor would make given the elicited γ. However, subjects 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, subjects who consume a positive fraction of asset A have to choose assets B and C in equal proportions[13] and ii) in weeks 7-10 no subject should invest in asset D as it is state-wise dominated, for example by a combination of assets A, B and C such as $(\frac{4}{10}, \frac{3}{10}, \frac{3}{10}, 0)$.[14]

After the tenth round, we conducted an additional questionnaire to elicit information on subjects' financial market experience, education, profession, and income. Finally, we randomly selected 100 subjects from all subjects 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

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

[12] For instance, if subjects had invested 20 GBP in asset A and 30 GBP in asset B, and state 2 realized, they would start the next period 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.

[13] The reason for this is that only risk-averse subjects 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.

[14] 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.

portfolio value of all subjects (19.22 GBP).

3 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. To 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) optimization in each period and is independent of wealth and of the investment horizon (Samuelson, 1969; Back, 2010).

Note first that our four-asset problem can 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 maximize 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 maximizers (and most non-expected utility maximizers) would not hold any asset D since it is dominated state-by-state e.g. by portfolio $2/5 * A + 3/5 * (\frac{1}{2}B, \frac{1}{2}C)$.

Thus, we consider an investor with wealth w who decides for one period 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 a return of r_1 and r_2 in these two states. The sure asset pays a return of s regardless 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

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

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. This is also the mix suggested by the robo-advisors.

3.1 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 maximizing 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 subjects, 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 GPB 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, 2017, 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 the appendix.

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.

4 Results

For the analysis of the results we shall follow closely our pre-analysis plan.[15] The first question outlined there concerns market participation, i.e. the question whether subjects are more willing to participate in our investment experiment if they are informed about existence of a robo-advisor. This question is important since robo-advisors have been praised for democratizing financial advise (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 compares the share of subjects who indicated their willingness to participate in the investment experiment by providing their email address.[16] Fisher exact tests (p-values > 0.63) show that there are no significant differences in the participation shares across treatments.

Result 1 *Offering a robo-advisor has no 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 4 and 5 in the appendix show summary statistics for observables across treatments. All statistics seem to be very close across treatments. In particular, for our most important variable, the coefficient of relative risk aversion γ, there are no significant differences across treatments among those who participated in the investment experiment (t-tests, $p > 0.31$).

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

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

[16] The corresponding figure for the first measure and a third (unregistered) measure, which counts how many subjects actually arrived in the investment stage are shown as Figures 8 and 9 in the appendix. None of the treatment differences are significant.

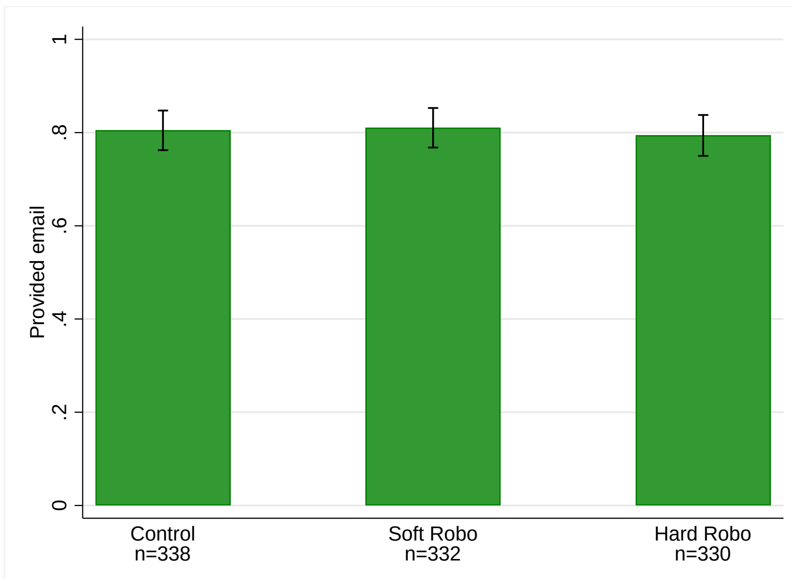

Figure 1: Shares of subjects who indicated their willingness to participate in the investment experiment by providing their email address.

Note: Shown are 95% confidence intervals.

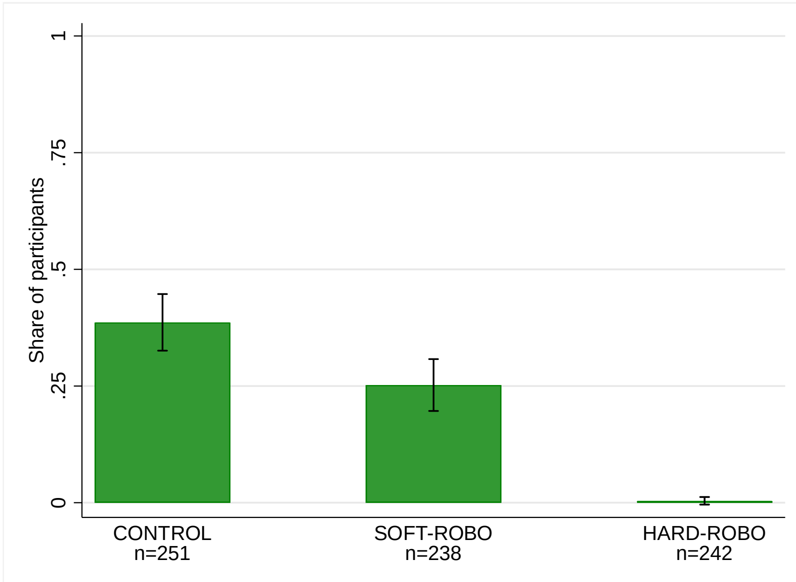

Figure 2: Share of investors who hold exclusively riskless asset A in week 10.
Note: Shown are 95% confidence intervals. 11 bankrupt investors are dropped.

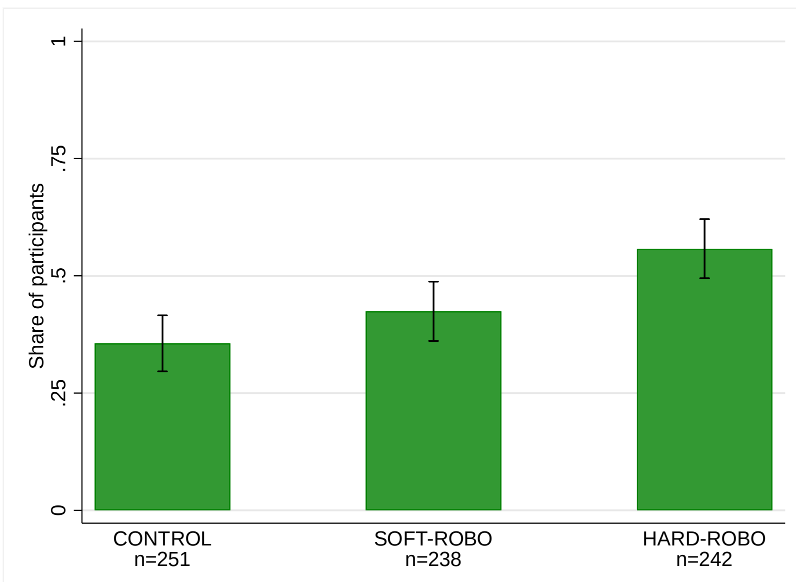

Figure 3: Share of investors who hold visited and clicked on the website in week 10
Note: Shown are 95% confidence intervals. 11 bankrupt investors are dropped.

study this question, we analyze the fraction of investors who invest *all* of their experimental wealth 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 in the last round. While 38.65% of investors do so in CONTROL, only 25.21% choose to do so in SOFT-ROBO, and a mere 0.41% do so in HARD-ROBO. All pair-wise 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 last round, dropping those investors who were bankrupt. Figure 3 reports the results. While 55.79% clicked in treatment HARD-ROBO and 42.44% in SOFT-ROBO, only 35.60% did 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 ($p = 0.137$).[17]

Result 2 *With robo-advisors there is significantly more continued participation in financial markets, 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 subjects?

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, in particular since asset D has a payoff structure that appeals to many given that it looks better than any of the other assets by themselves. 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

[17] 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.

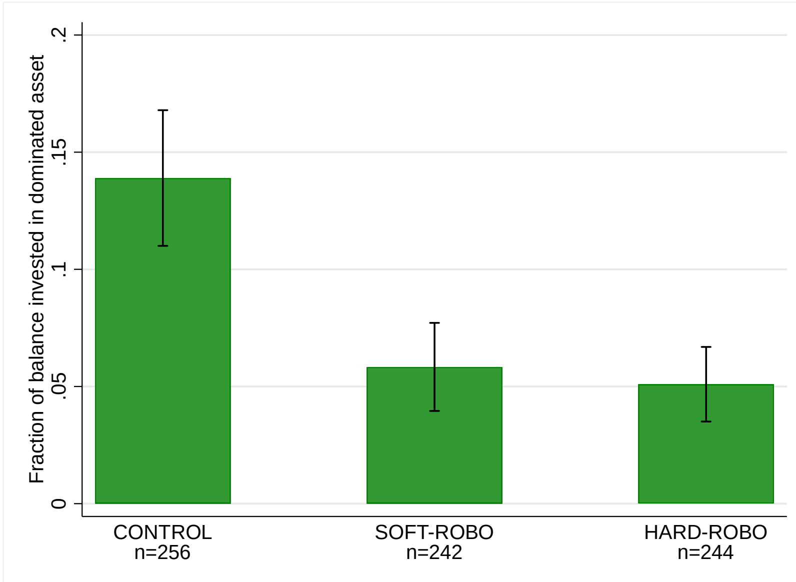

Figure 4: Fraction of balances invested in dominated asset D.
Note: Shown are 95% confidence intervals.

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. 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."

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

One of the most important tasks of actual robo-advisers and which most banks advertise is rebalancing. After the returns of the risky assets are drawn at the end of each period, investors have a new default portfolio, which will be carried over into the next period 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 CE of the default portfolio ($CE_t^{default}$). Figure 5

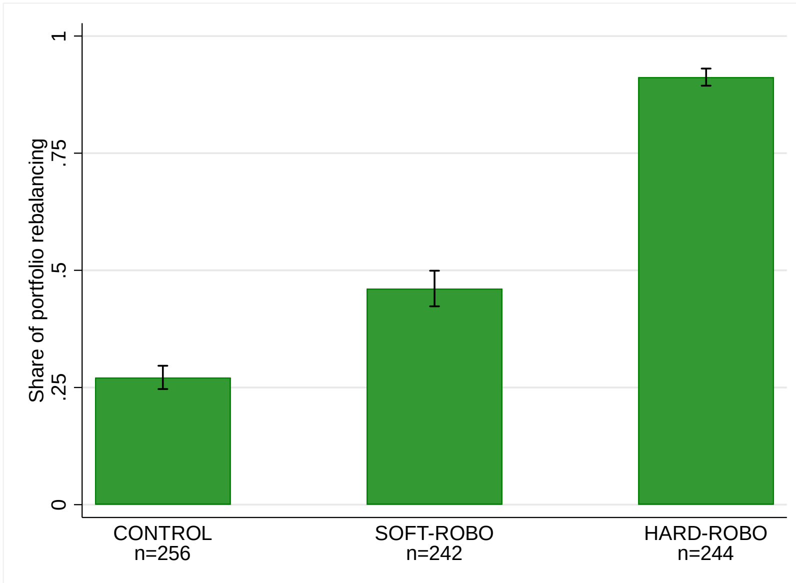

Figure 5: How often are portfolios rebalanced in the right direction?

shows the share of portfolios which are rebalanced in the right direction. Since the Hard robo rebalances automatically, it is not surprising that this share is almost 1, although in about 9% of periods investors override the robo. In SOFT-ROBO, rebalancing happens in only 46% of cases and in CONTROL in only 27% of cases. All pairwise differences are significant (using each investor as one independent observation, t-test, $p < 0.0001$). 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 login every week.

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

Standard portfolio theory requires that risks are diversified as much as possible. 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

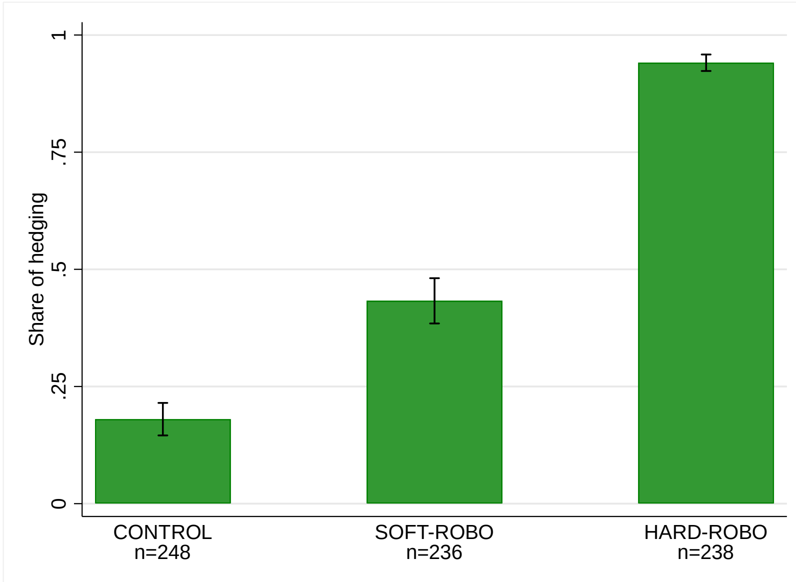

Figure 6: How often do investors hedge uncertainty?
Note: Shown are 95% confidence intervals. If an investors holds any asset A, then a correct hedge requires that the weights of assets B and C do not differ by more than 10%. 20 investors never held any asset A.

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 subjects. 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.[18]

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. 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 advise (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

[18] 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.

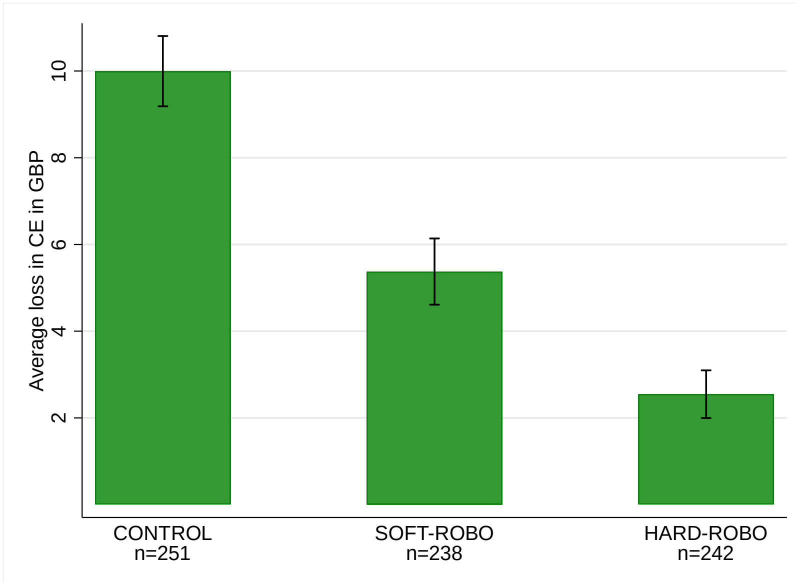

Figure 7: Average loss in certainty equivalent over 10 weeks
Note: Shown are 95% confidence intervals. We dropped 11 subjcts who went bankrupt.

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, $p < 0.0001$).[19]

Result 6 *Using a Robo can avoid substantial losses in CE, with Hard robos again significantly outperforming Soft robos.*

This concludes the analysis as registered in our pre-analysis plan. However, given that we have seen that many investors do not follow the robo's recommendations, as an exploratory exercise, we would like to understand better why this happens. For this we count how often investors choose a portfolio that meets for each asset the robo's recommended amount plus or minus 10%. In SOFT-ROBO, investors follow the robo's advice in only about 25% of cases. This increases to 44% when investors were actually visiting our website in the given week. In HARD-ROBO, investors follow the robo's advice in 82% of cases

[19] In the experiment 11 subjects went bankrupt due to rounding of the balance. To be conservative, we dropped them for this analysis since their $\overline{CE}_1$ would be zero and only 2 of the 11 subjects come from HARD-ROBO.

but this goes down to 71% when they actually visited the webpage (recall that the robo's advice is implemented by default if investors do not visit the webpage). 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 run a linear probability model (see Table 3) for the probability of following advice (within a 10% band). We suspect that investors experience regret if they followed the robo's advice in the previous week and then were disappointed by the outcome. However, it is not straightforward to define what constitutes regret as we cannot observe what subjects would have chosen without a robo. Would they have invested less or more in the risky assets on their own? Since we cannot know this, we resort to using the occurrence of State 3, in which they are sure to lose with all risky assets, as a proxy for a disappointing outcome. In the regressions we use only subject-week combinations in which (1) subjects visited the website in this week (because if they did not, they cannot override the advice) and (2) subjects held some amount of risky assets in the previous week (because if they did not, they cannot experience regret).

Table 3 shows that investors are more likely to follow advice in the current week when they already followed advice in the previous week, which is not surprising. In treatment HARD-ROBO, the occurrence of State 3 in the previous week increases the probability of following advice by about 20 percentage points *unless* the investors followed the advice in the previous week and were disappointed with it due to the occurrence of State 3. These effects are much smaller and not significant in treatment SOFT-ROBO. Adding additional controls to the regressions does not change much. The existence of the dominated asset lowers the probability of following advice, probably because some investors are convinced of the dominated asset despite of the advice. Higher balance also leads to a lower probability of following advice. Finally, in HARD-ROBO, females are about 6% less likely to follow advice.

5 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 found that investors' utility was increased substantially by having access to a robo-advisor. The benefits were larger for the investors in the HARD-ROBO treatment, where robo-advice was implemented automatically, than for those in treatment SOFT-ROBO, where investors simply received recommendations and had to implement them on their own. This is in line with the evidence on the power of default settings (see e.g. Choi et al., 2004).

In our experiment the benefits of robo-advisors materialize because of their ability to (1) select and tailor the initial portfolio according to risk preferences, (2) rebalance portfolios after each week, (3) avoid a dominated asset, and (4) use mixtures of assets to hedge risk over some states. Many real world robo-advisors are designed to do just that (Vanguard, 2023). Yet arguably, how important these features are in a real world setting may depend on the particular nature of the financial market under consideration.[20]

Existing empirical studies probably often underestimate the benefits of robo-advisors because they compare expected or realized returns. Robo-advisors are designed to maximize expected utility of investors and only risk-neutral investors would be interested in maximizing expected returns. In this study we compared the difference in certainty equivalents of the portfolios with and without robo advice and found that HARD-ROBOs helped 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.[21]

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

[20] For example, the existence of dominated assets or hedging opportunities may be less transparent in some real markets while they are important in other markets. For example, financial instruments with excessive fees would correspond to dominated assets or hedging opportunities may arise in the form of internationally diversified ETFs.

[21] Of course, these quantitative results may not be readily portable to real life financial markets as they reflect the parameterization chosen in our experiment. Notwithstanding, our results point towards substantial welfare gains, which may be overlooked when considering expected returns rather than expected utility.

Table 3: Linear probability model: Do investors follow the robo's advice?

	SOFT-ROBO		HARD-ROBO	
followed advice in prev. week	0.413^{***}	0.411^{***}	0.521^{***}	0.497^{***}
	(0.039)	(0.039)	(0.044)	(0.045)
State 3 last week	0.050	0.028	0.223^{***}	0.175^{***}
	(0.039)	(0.041)	(0.060)	(0.063)
followed_advice × State3 last week	−0.085	−0.087	-0.286^{***}	-0.271^{***}
	(0.061)	(0.062)	(0.064)	(0.065)
week with dominated asset		-0.151^{***}		-0.103^{**}
		(0.057)		(0.044)
week		0.014		0.002
		(0.011)		(0.008)
balance		-0.002^{*}		-0.002^{**}
		(0.001)		(0.001)
γ		−0.010		−0.012
		(0.019)		(0.014)
financial_literacy		0.034		−0.024
		(0.039)		(0.029)
age		-0.003^{**}		0.000
		(0.002)		(0.001)
female		−0.037		-0.059^{**}
		(0.040)		(0.029)
constant		0.514^{***}		0.509^{***}
		(0.103)		(0.089)
R^2	0.142	0.167	0.187	0.201
Observations	1039	1015	1239	1231

Note: Follow advice is counted as 1 if amount invested in each asset is within 10% of recommended amount. Subject-weeks are included only if (1) subjects visited the website in this week and (2) subjects held some amount of risky assets in the previous week. Standard errors (in parentheses) are clustered on the subject level. $^{*}p < 0.10$, $^{**}p < 0.05$, $p < 0.01$.

References

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

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

Barberis, N. and Thaler, R. (2003). A survey of behavioral finance. *Handbook of the Economics of Finance*, 1:1053–1128.

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.

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

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 *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. *The Palgrave Handbook of Technological Finance*, pages 725–749.

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. 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. *Handbook of Experimental Economics Results*, 1:1061–1073.

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

Gaudeul, A. and Giannetti, C. (2023). Trade-offs in the design of financial algorithms. *Available at SSRN 4432707*.

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.

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.

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.

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.

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*.

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

Rossi, A. G. and Utkus, S. P. (2020). Who benefits from robo-advising? Evidence from machine learning. *mimeo*.

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

Statista (2023). Assets under management of robo-advisors worldwide from 2018 to 2027. https://www.statista.com/forecasts/1262614/robo-advisors-managing-assets-worldwide. Accessed on July 4, 2023.

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*.

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

Appendix A: Proofs

Proposition 1 *The CE of the dynamic portfolio problem with CRRA utility and returns that follow an i.i.d. process is equal to the product of the CEs of the per period CEs*

Note that the proposition assumes that returns follow an i.i.d. process. Thus, to apply it to our experimental design, we have to consider weeks 1-3 and 4-10 separately.

Proof Consider the last 2 weeks. In week T, choose portfolio weight f with corresponding payoffs $F_1 = f(1 + r_1) + (1 - f)(1 + s)$ and $F_2 = f(1 + r_2) + (1 - f)(1 + s)$. In week $T - 1$ a possibly different weight g is chosen such that $G_1 = g(1 + r_1) + (1 - g)(1 + s)$ and $G_2 = g(1 + r_2) + (1 - g)(1 + s)$. Let L_t denote the lottery induced by the portfolio choice in week t.

Expected utilities are then $U(L_T) = pu(F_1) + (1 - p)u(F_2)$ and $U(L_{T-1}) = pu(G_1) + (1 - p)u(G_2)$.

$$\begin{aligned}
& U(L_T)U(L_{T-1}) \\
= \; & p^2u(F_1)u(G_1) + p(1-p)u(F_1)u(G_2) + p(1-p)u(G_1)u(F_2) + (1-p)^2u(F_2)u(G_2) \\
& U(L_T \otimes L_{T-1}) \\
= \; & p^2u(F_1G_1) + p(1-p)u(F_1G_2) + p(1-p)u(G_1F_2) + (1-p)^2u(F_2G_2),
\end{aligned}$$

where $\otimes$ denotes the compound lottery.

Note that all terms have same probabilities. And for CRRA utility we have

$$u(F_i)u(G_j) = \frac{F_i^{1-\gamma}}{1-\gamma}\frac{G_j^{1-\gamma}}{1-\gamma} = \frac{(F_iG_j)^{1-\gamma}}{(1-\gamma)^2} = \frac{1}{1-\gamma}u(F_iG_j).$$

Hence,

$$(1-\gamma)\, U(L_T)U(L_{T-1}) = U(L_T \otimes L_{T-1}).$$

Thus

$$\begin{aligned}
(1-\gamma)\, U(CE_T)U(CE_{T-1}) &= U(\overline{CE}_{T-1}) \\
(1-\gamma)\, \frac{CE_T^{1-\gamma}}{1-\gamma}\frac{CE_{T-1}^{1-\gamma}}{1-\gamma} &= \frac{\overline{CE}^{1-\gamma}}{1-\gamma} \\
CE_T CE_{T-1} &= \overline{CE}_{T-1}.
\end{aligned}$$

Working backwards we can now conclude that

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

Appendix B: Additional figures and tables

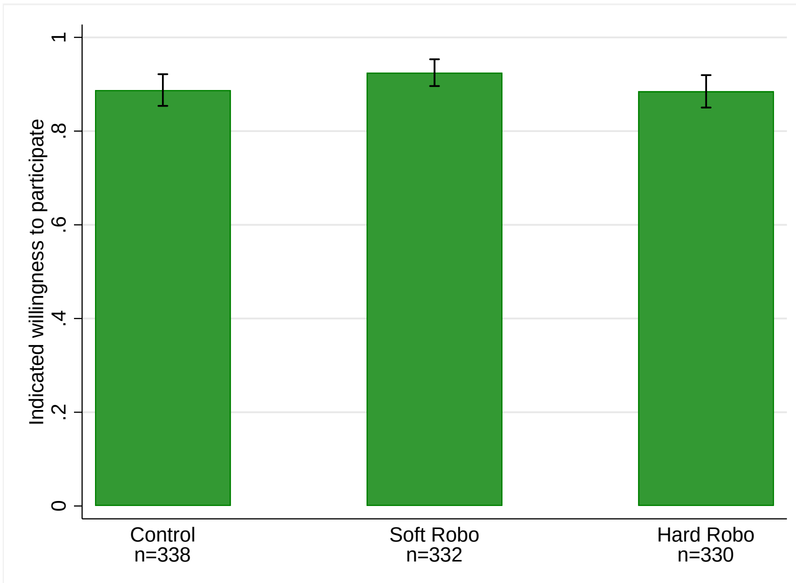

Figure 8: Shares of subjects who said they wanted to participate in the investment experiment.

Note: Shown are 95% confidence intervals.

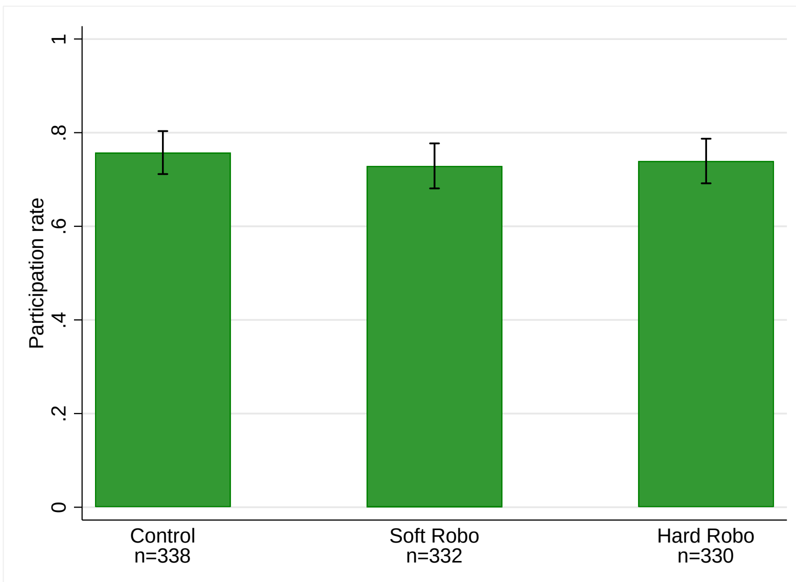

Figure 9: Shares of subjects who actually arrived at the investment stage of the experiment.
Note: Shown are 95% confidence intervals.

Table 4: Summary statistics, all subjects

	Mean	SD	Min	Max	N
CONTROL					
risk_category	2.24	1.11	1	5	338
investment_experience	0.46	0.50	0	1	338
know_robo	0.34	0.47	0	1	338
used_robo	0.06	0.24	0	1	338
age	39.00	12.63	18	74	337
female	0.51	0.50	0	1	330
financial_market_experience	1.61	0.72	1	3	122
profession_category	5.42	2.64	1	8	122
income_category	3.04	1.56	1	6	122
financial_literacy	0.49	0.50	0	1	338
SOFT-ROBO					
risk_category	2.38	1.14	1	5	332
investment_experience	0.47	0.50	0	1	332
know_robo	0.60	0.49	0	1	332
used_robo	0.06	0.23	0	1	332
age	39.50	12.90	18	74	330
female	0.50	0.50	0	1	331
financial_market_experience	1.76	0.75	1	4	127
profession_category	5.31	2.61	1	8	127
income_category	3.27	1.61	1	6	127
financial_literacy	0.48	0.50	0	1	332
HARD-ROBO					
risk_category	2.24	1.08	1	5	330
investment_experience	0.44	0.50	0	1	330
know_robo	0.58	0.49	0	1	330
used_robo	0.05	0.22	0	1	330
age	40.63	13.62	19	87	327
female	0.49	0.50	0	1	329
financial_market_experience	1.75	0.75	1	4	157
profession_category	5.02	2.76	1	8	157
income_category	2.97	1.59	1	6	157
financial_literacy	0.54	0.50	0	1	330

Note: We report mean, standard deviation, minimum, maximum, and the number of observations. We receive information on gender and age from Prolific, while all other variables are from our questionnaires.

Table 5: Summary statistics, subjects who reached investment game

	Mean	SD	Min	Max	N
CONTROL					
risk_category	2.33	1.08	1	5	256
investment_experience	0.50	0.50	0	1	256
know_robo	0.36	0.48	0	1	256
used_robo	0.05	0.22	0	1	256
age	38.94	12.56	18	73	255
female	0.50	0.50	0	1	253
financial_market_experience	1.61	0.72	1	3	122
profession_category	5.42	2.64	1	8	122
income_category	3.04	1.56	1	6	122
financial_literacy	0.54	0.50	0	1	256
SOFT-ROBO					
risk_category	2.40	1.14	1	5	242
investment_experience	0.51	0.50	0	1	242
know_robo	0.68	0.47	0	1	242
used_robo	0.06	0.23	0	1	242
age	40.17	12.18	18	73	240
female	0.47	0.50	0	1	241
financial_market_experience	1.76	0.75	1	4	127
profession_category	5.31	2.61	1	8	127
income_category	3.27	1.61	1	6	127
financial_literacy	0.51	0.50	0	1	242
HARD-ROBO					
risk_category	2.30	1.06	1	5	244
investment_experience	0.50	0.50	0	1	244
know_robo	0.63	0.48	0	1	244
used_robo	0.06	0.23	0	1	244
age	40.15	13.21	19	72	242
female	0.48	0.50	0	1	244
financial_market_experience	1.75	0.75	1	4	157
profession_category	5.02	2.76	1	8	157
income_category	2.97	1.59	1	6	157
financial_literacy	0.56	0.50	0	1	244

Appendix C: Experimental Instructions & Details

The first stage of our experiment was designed to recruit subjects for the second stage, our investment game. During the recruitment survey subjects had to go through a couple of tasks in fixed order. Subsection 5.1.1 describes our numeracy task. Depending on the treatment, we provided our subjects with a description of the follow-up investment game and whether they would have access to a robo-advisor. We include these in subsection 5.1.2. In subsection 5.1.3, we describe the incentivized quiz regarding the follow-up experiment. Subsection 5.1.4 describes the task we used to elicit risk preferences. We asked subjects whether they have investment experience and have knowledge about robo-advisors, see subsection 5.1.5. In subsection 5.1.6 we show how we invited subjects to join the follow-up experiment. When joining the investment game stage, subjects had the opportunity to sign up for the experiment by providing their informed consent and their e-mail address, see subsection 5.2.1. Subsequently, they received detailed instructions on the investment game and (treatment dependent) information about their robo-advisor. To ensure that subjects understood the instructions, they had to pass comprehension quiz questions before being able to proceed to the investment game. Subsection 5.2.2 provides details on both the instructions and the corresponding quiz questions. In subsection 5.2.3, we include and describe (fictional) screens of the investment game stage. Finally, we list the questions of our debriefing survey in subsection 5.2.4. The following table provides a summary of the stages in our experiment.

		Stage
1		Recruitment survey (1000 subjects, about 15 minutes)
	1.1	Incentivized numeracy test (BNT)
	1.2	Description of follow-up experiment
	1.3	Incentivized quiz about the follow-up experiment
	1.4	Incentivized risk elicitation task (Binswanger)
	1.5	Investment and robo experience survey
	1.6	Offer to join the follow-up experiment
2		Investment game (10 weeks)
	2.1	Informed consent and email registration
	2.2	Instructions and quiz
	2.3	Investment game
	2.4	Demographics survey

5.1 Recruitment stage

5.1.1 Numeracy task

We implement the single item version of the Berlin Numeracy Test (Cokely et al. 2012). According to the authors, this version is designed for situations where time is limited and allows to estimate median splits. Those who answer the question right are estimated to belong to the top half of educated subjects. We offered subjects a bonus payment of 0.50 GBP if they answer correctly, and asked them not to use a calculator.

> Out of 1,000 people in a small town 500 are members of a choir. Out of these 500 members in the choir 100 are men. Out of the 500 inhabitants that are not in the choir 300 are men. What is the probability that a randomly drawn man is a member of the choir? Please indicate the probability in percent. *Correct answer: 25*

5.1.2 Description of follow-up experiment

During the recruitment stage, we inform subjects that they will have the chance to sign up for a follow-up experiment. They are also informed that careful reading can earn them some bonus money by answering a few quiz questions. The first screen shows some general information:

> In this stage of the experiment we will inform you about a **follow-up experiment** that we are running. Later on, you may decide whether you want to join the follow-up experiment. Note that your participation is optional and your payment in the current experiment is unaffected by the participation decision. However, you can **earn extra points** by correctly answering questions regarding the details of the follow-up experiment.
>
> In the follow-up experiment, participants start with **50 GBP** and can make an **investment decision once per week** over the time span of **10 weeks**. In total, we invite **1,000 participants** to participate in our study. At the end, we will randomly draw **100 participants** out of those who choose to participate in the follow-up experiment and **pay** them their **actual portfolio value in GBP**.
>
> On the next page, you will find more details about the follow-up experiment. **Please read them carefully**. Subsequently, you can earn bonus money for

correctly answering questions that refer to the details of the follow-up experiment. In particular, you will be presented with **three multiple choice questions**. Each question has **one correct answer**. For each question that you answer correctly you will earn **100 points**. You will learn about your correct answers in the payment summary at the end of this experiment.

The second screen describes the follow-up experiment in more detail:

In the follow-up experiment, participants start with **50 GBP** and can make an **investment decision once per week** over the time span of **10 weeks**. In each week, a **set of assets** is available to invest into. You may decide how to **invest your current holdings** across the available assets. Each of the assets pays GBP according to your investment and a **randomly drawn state**. Each asset has **three different states** of realization, and each of the three states is **equally likely** to be realized.

Here is an example (involving just a single asset):

	State 1	State 2	State 3
Asset A	1.50	1.00	0.50

The table shows how the value of your portfolio changes for every GBP you invest in Asset A: If State 1 releases, the value of your portfolio increases by 50%, stays the same if State 2 realises, and reduces by 50% if State 3 realises.

Suppose you invest your initial 50 GBP in Asset A.
Then, with probability 1/3 the value of your portfolio will be 50 x 1.50 = 75 GBP in the next week.
Likewise, with probability 1/3 the value of your portfolio will be 50 x 1.00 = 50 GBP in the next week.
Finally, with probability 1/3 the value of your portfolio will be 50 x 0.50 = 25 GBP in the next week.
In each week, you have to **fully invest** your funds across the available assets.

NOTE: The following paragraph was exclusively included in the SOFT-ROBO treatment

You will receive advice from a **robo-advisor** regarding your investment decisions. The robo-advisor will **account** for your **attitude towards risks** and will recommend how to split up your current holdings across the available assets. Your risk attitude will be measured using a **lottery task**. Thus, in order to allow the robo-advisor to make the **best recommendation** for you, you should answer the questions in the lottery task as **sincerely** as possible.

However, the recommendation of the robo-advisor during the investment game will **not be binding**. In each week, you will need to **specify your investment manually**, and it is up to you whether you follow the recommendation of the robo-advisor or deviate from it. If you do not specify an investment strategy in a given week, a default investment will be implemented.

NOTE: The following paragraph was exclusively included in the HARD-ROBO treatment

You will receive advice from a **robo-advisor** regarding your investment decisions. The robo-advisor will **account** for your **attitude towards risks** and will recommend how to split up your current holdings across the available assets. Your risk attitude will be measured using a **lottery task**. Thus, in order to allow the robo-advisor to make the **best recommendation** for you, you should answer the questions in the lottery task as **sincerely** as possible.

In each week, the **robo-advisor** will **invest your current holdings** for you. However, in each week, you have the option to **override** the investment proposed by the robo-advisor and thus deviate from it. If you do not override the robo investment in a given week, your current holdings will be invested according to the strategy proposed by the robo-advisor.

Prior to the investment game, we will ask for your e-mail to send you a **weekly reminder** of the ongoing experiment. We will also use your e-mail to coordinate payments at the end of the investment game and will **delete your email address** at the end of the experiment. In total, we invite **1,000 participants** to participate in our study. For compensation, we will **randomly** draw **100 participants** out of those who choose to participate in the follow-up experiment and pay them their **full portfolio value** in GBP. Payments will be implemented via direct bank transfer or Amazon vouchers depending on your preference.

5.1.3 Quiz in recruitment stage

Before providing details on the follow-up experiment we announced quiz questions which may lead to a bonus payment of up to 300 ECU ($\approx$ 0.38 GBP). Our main intention was to motivate attention and careful reading. The questions cover the general setup of the experiment and are not treatment dependent.[22]

- What are the starting funds in the investment game of the follow-up experiment? (10 GBP; 20 GBP; 50 GBP - *Correct answer: 50 GBP*)
- What is the time frame of the follow-up experiment? (10 minutes, with one decision every minute; 10 weeks, with one decision every week; 10 months, with one decision every month - *Correct answer: 10 weeks, with one decision every week*)
- Who will get paid for participation in the follow-up experiment? (100 randomly drawn participants receive their full portfolio value at the end of the experiment; Every participant receives their full portfolio value at the end of the experiment; Every participant receives 10 GBP - *Correct answer: 100 randomly drawn participants receive their full portfolio value at the end of the experiment*)

5.1.4 Risk preference elicitation

We elicit risk preferences of our participants using the Binswanger task (cf. Binswanger 1980; Eckel and Grossman 2008). Table 1 depicts the lotteries and corresponding rewards in ECU. Participants in treatments SOFT-ROBO and HARD-ROBO were informed beforehand that their decision would calibrate the robo-advisor. They were not informed about the technical implementation, i.e. that we assign them a risk aversion parameter representative of their choice which would be used by the robo to maximize a CRRA utility function with that parameter. Unfortunately, we made a mistake when re-scaling the outcomes of the safe lottery and offered 760 ECU instead of 740 ECU. This means that slightly less risk averse subjects may have still have chosen the safe lottery. However, as this category also includes the most extreme risk averse subjects, we still feel confident to have imputed a representative risk preference parameter for that category. Furthermore,

[22] For quiz questions regarding the details of the investment game and the robo advisor, see subsection 5.2.2.

the optimal portfolio choice is (fortunately) inelastic with respect to small changes in γ, in particular for very risk averse subjects who chose the first or second lottery.

5.1.5 Survey about investment and robo experience

We asked participants to answer three unincentivized questions about their investment experience and whether they had heard about or used robo-advisors previously.

> The next stage of the experiment is a **short survey** of three questions. Please answer these questions truthfully, as they are very important for our research.

- Have you ever invested money in the stock market?
  (Yes; No)
- Do you know what a robo-advisor is?
  (Yes; No)
- Have you ever used a robo-advisor?
  (Yes; No)

5.1.6 Offer to join the follow-up experiment

Finally, we asked our participants whether they are interested in registering for the follow-up experiment. For those who chose "Yes" we created a personal link to refer them to the registration of the investment game stage.

> Thank you for completing the tasks of this experiment. You can now decide whether you want to participate in the follow-up experiment.

- Would you like to participate in the follow-up experiment?
  (Yes; No)

5.2 Investment game stage

5.2.1 Registration for the follow-up experiment

When joining our follow-up experiment participants first had to provide their informed consent and an e-mail address.

Thank you for joining our follow-up experiment. Before we start, we would like to ask for your **e-mail address.**

By providing us with your e-mail address you consent to us sending you regular reminders about your portfolio. We will delete your e-mail address at the end of the experiment and you will receive no further e-mails from us. We will use it **exclusively** in relation to this experiment, in particular to send you a **reminder** once a week about the **decision** that you can make, and to provide your **personal link** that takes you back to the experiment. We will also use your e-mail to contact you regarding **payment of the experiment**. Your e-mail address will be stored **separately** and not be connected to any other data. We will only use your email address for the purposes mentioned above and will delete it after the experiment is completed. Please make sure that you receive our e-mails, which could mean that you need to check your spam folder. In case you are in doubt about receiving our e-mails, please contact XXX.

By providing your e-mail address you also confirm that you are not an employee of the University of Essex who is participating during their normal employment hours.

5.2.2 Investment game instructions and quiz

After successful registration participants received detailed instructions on the investment game.

This experiment is about an **investment game**. You start with **50 GBP** and can make an **investment decision once per week** over the time span of **10 weeks**. In total, we invite **1,000 participants** to participate in our study. At the end, we will **randomly** draw **100 participants** out of those who choose to participate in the follow-up experiment and **pay** them their **actual portfolio value in GBP**. Note that in order to **qualify for the payment lottery**, you need to **complete the instructions stage** before Sunday, May 22th, 23:59 BST. If you will be drawn for payment, you will receive exactly the amount of money that you hold at the end of the experiment. The amount of money you hold at the end of the experiment will depend on your **investment decisions** and **random draws**.

You will find **detailed instructions** on the **next page**. Please **read** these instructions **carefully**. Before you start the game, we will ask you a few **questions** about how things work to ensure that the rules are communicated clearly. You can only **advance to the game** once you **answer all questions correctly**.

In this experiment, participants start with **50 GBP** and can make an **investment decision once per week** over the time span of **10 weeks**.

In each week, a **set of assets** is available to invest into. Note that the set of available assets may change over the course of the experiment. Whenever the set of available assets has changed, you will receive a notification about it.

Your task during the experiment is to decide how to **invest your current holdings** across the available assets. Each of the assets pays GBP according to your investment and a **randomly drawn state**. Each asset has **three different states** of realization, and each of the three states is **equally likely** to be realized.

Here is an example (involving just a single asset):

	State 1	State 2	State 3
Asset A	1.50	1.00	0.50

The table shows how the value of your portfolio changes for every GBP you invest in Asset A:

If State 1 realises, the value of your portfolio increases by 50%, stays the same if State 2 realises, and reduces by 50% if State 3 realises.

Suppose you invest your initial 50 GBP in Asset A.
Then, with probability 1/3 the value of your portfolio will be 50 x 1.50 = 75 GBP in the next week.
Likewise, with probability 1/3 the value of your portfolio will be 50 x 1.00 = 50 GBP in the next week.
Finally, with probability 1/3 the value of your portfolio will be 50 x 0.50 = 25 GBP in the next week.
In each week, you have to **fully invest** your funds across the available assets.

NOTE: The following paragraph was exclusively included in the CONTROL treatment

When submitting an investment strategy, you will see a **summary of potential outcomes** depending on the different states **before your strategy**

becomes binding. If you do not specify an investment strategy in a given week, your current holdings will remain invested as they were at the beginning of the week.

NOTE: The following paragraph was exclusively included in the SOFT-ROBO treatment

You will receive advice from a **robo-advisor** regarding your investment decisions. The robo-advisor will **account** for your **attitude towards risks** and will recommend how to split up your current holdings across the available assets. Your risk attitude was measured during the preceding experiment and used to **calibrate the robo-advisor** accordingly. However, the recommendation of the robo-advisor during the investment game will **not be binding**. In each week, you will need to **specify your investment manually**, and it is up to you whether you follow the recommendation of the robo-advisor or deviate from it. When submitting an investment strategy, you will see a **summary of potential outcomes** depending on the different states **before your strategy becomes binding**. If you do not specify an investment strategy in a given week, your current holdings will remain invested as they were at the beginning of the week.

NOTE: The following paragraph was exclusively included in the HARD-ROBO treatment

You will receive advice from a **robo-advisor** regarding your investment decisions. The robo-advisor will **account** for your **attitude towards risks** and will recommend how to split up your current holdings across the available assets. Your risk attitude was measured during the preceding experiment and used to **calibrate the robo-advisor** accordingly.

In each week, the **robo-advisor** will **invest your current holdings** for you. However, in each week, you have the option to **override** the investment proposed by the robo-advisor and thus deviate from it. If you do not override the robo investment in a given week, your current holdings will be invested according to the strategy proposed by the robo-advisor.

Prior to the investment game, we asked for your e-mail, which we will use to send you a **weekly reminder** of the ongoing experiment. In those e-mails, we will also include your **personal link** to get back to the experiment. Please do not share this link with anyone. It is also possible to **bookmark** the webpage. Finally, we will use your e-mail to coordinate payments at the end of the investment game and will **delete your e-mail address** at the end of the experiment. In total, we invite **1,000 participants** to participate in our study. For compensation, we will **randomly** draw **100 participants** out of those who choose to participate in the follow-up experiment and pay them their **full portfolio value** in GBP. Payments will be implemented via direct bank transfer or Amazon vouchers depending on your preference.

In order to make sure that participants had understood the key aspects of the investment game, we had them answer a few quiz questions. Only after answering all questions correctly they could move on to the investment game. In case they answered one or more questions incorrectly they were returned to the instruction page.

- What influences the value of your portfolio at the end of each week?
  (Your investment; A random draw; A combination of your investment and a random draw - *Correct answer: A combination of your investment and a random draw*)

- Does the set of available assets change over the course of the experiment?
  (Yes, the set of available assets may change, and you will receive notifications accordingly; No, the set of available assets is constant over the course of the experiment - *Correct answer: Yes, the set of available assets may change, and you will receive notifications accordingly*)

- (Treatment CONTROL) What happens if you do not submit an investment strategy in a given week?
  (Your current holdings will transfer to the next week without being invested; Your current holdings will remain invested as they were at the beginning of the week; You will lose your entire holdings and be disqualified from the experiment - *Correct answer: Your current holdings will remain invested as they were at the beginning of the week*)

- (Treatment SOFT-ROBO and HARD-ROBO) What happens if you do not submit an investment strategy in a given week?

(Your current holdings will transfer to the next week without being invested; Your current holdings will remain invested as they were at the beginning of the week; Your current holdings will be invested according to the robo-advisor recommendation; You will lose your entire holdings and be disqualified from the experiment - *Correct answer in SOFT-ROBO: Your current holdings will remain invested as they were at the beginning of the week; in HARD-ROBO: Your current holdings will be invested according to the robo-advisor recommendation*)

5.2.3 Investment game

In this section we provide fictional example screens of the investment game stage. All treatments share the same general information at the top of the decision page (Figure 10). Here, participants can see how much time they have left to submit a decision, important notifications, and the evolution of their portfolio over time. Figure 11 shows the information participants receive about the asset scenario of the current week. In treatment HARD-ROBO, they can then see how the advisor would invest holdings across the available assets, and accept the recommendation if they wish to do so. In treatment SOFT-ROBO, participants see both their current investment and the investment recommendation by the robo (Figure 12). Participants in treatment CONTROL only receive information about the asset scenario and about their current investment. Irrespective of the treatment, participants can submit their own investment strategies (either to overwrite the robo recommendation in HARD-ROBO, or to change their holdings from the beginning of the week in treatments SOFT-ROBO and CONTROL). Figure 13 shows the interface to submit an individual investment plan. After specifying the investment and clicking the (treatment-dependent) confirmation button, participants see a popup window that informs them about potential consequences by providing state-dependent scenarios. At this point, participants may also decide to not submit their investment plan and return to the decision screen instead.

Figure 10: Fictional example screen of the investment game: General information

▼ Available assets:

	State 1	State 2	State 3
Asset A	0.90	0.90	0.90
Asset B	2.80	0.10	0.10
Asset C	0.10	2.80	0.10
Asset D	1.20	1.20	0.36

The Table shows how the value of your portfolio changes for every GBP you invest.

Here is an example:
Suppose you invest 20 GBP in asset A and 30 GBP in asset B.

If the random mechanism picks state 1 for implementation, your investment in asset A will turn into holdings of 20 x 0.90 = 18 GBP in the next round, and your investment in asset B will turn into holdings of 30 x 2.80 = 84 GBP.

If the random mechanism picks state 2 for implementation, your investment in asset A will turn into holdings of 20 x 0.90 = 18 GBP in the next round, and your investment in asset B will turn into holdings of 30 x 0.10 = 3 GBP.

If the random mechanism picks state 3 for implementation, your investment in asset A will turn into holdings of 20 x 0.90 = 18 GBP in the next round, and your investment in asset B will turn into holdings of 30 x 0.10 = 3 GBP.

Note: Whenever you override the robo advisor recommendation, you will see a summary of potential outcomes depending on the different states before your strategy becomes binding.

Note: The robo advisor has invested your current holdings. If you want to deviate from this investment, you need to override it.

▼ Your investment strategy set by the robo advisor:

Asset	Investment (in GBP)
Asset A	12.69
Asset B	2.05
Asset C	2.05
Asset D	0.00

Accept robo recommendation

Figure 11: Fictional example screen of the investment game: Asset information (HARD-ROBO)

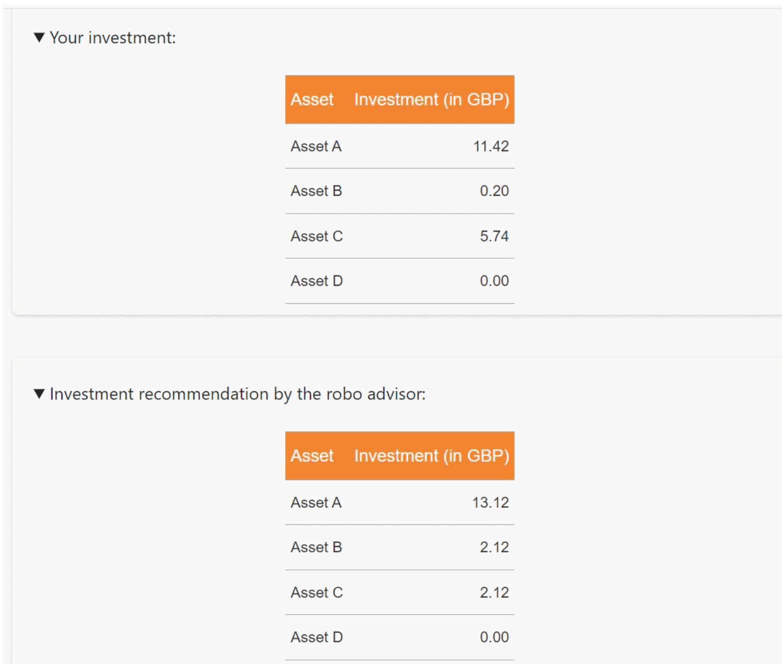
▼ Your investment:

Asset	Investment (in GBP)
Asset A	11.42
Asset B	0.20
Asset C	5.74
Asset D	0.00

▼ Investment recommendation by the robo advisor:

Asset	Investment (in GBP)
Asset A	13.12
Asset B	2.12
Asset C	2.12
Asset D	0.00

Figure 12: Fictional example screen of the investment game (SOFT-ROBO)

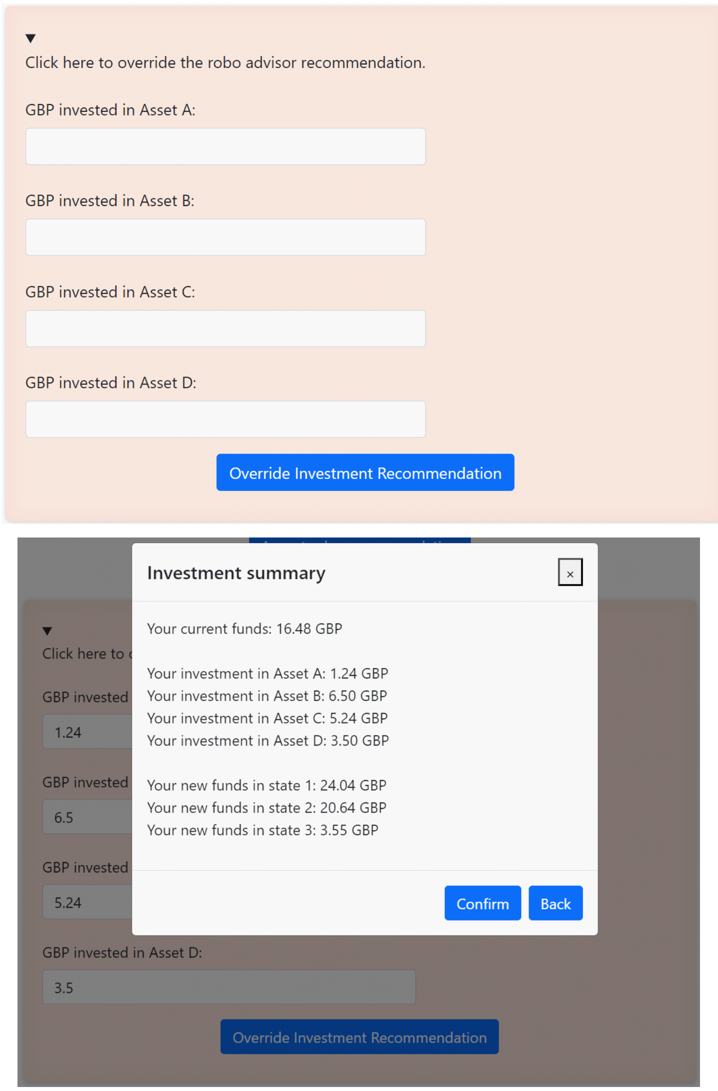

Figure 13: Fictional example screen of the investment game: Investment decision

5.2.4 Debriefing survey

After the end of the investment game we asked participants to finalize our survey by sharing some of their demographics. Note that we already had received information on their age and gender through Prolific.

- Do you have experience in financial markets and investment decisions?
  (No experience; Rather little; Above average; Very experienced)
- What is your highest level of education?
  (No formal education; Primary education; Secondary education (High school); Bachelor degree; Master degree; PhD or higher)
- What is your profession?
  (No profession; Arts and Entertainment; Business; Industrial and Manufacturing; Law Enforcement and Armed Forces; Science and Technology; Healthcare and Medicine; Other)
- What is your (approximate) annual income (in GBP)?
  (below $10,000$; between $10,000$ and $20,000$; between $20,000$ and $30,000$; between $30,000$ and $40,000$; between $40,000$ and $50,000$; higher than $50,000$)