Reichman University
Arison ESG Center
Environmental. Social. Governance

Arison ESG Center Working Papers No: 6

The Asymmetry in Responsible Investing Preferences and Beliefs

Jacquelyn Humphrey
Shimon Kogan
Jacob Sagi
Laura Starks

February 2024

Arison ESG Center, Arison Business School, Reichman University,
HaUniversita Street 8, Herzliya, Israel, esg@runi.ac.il

The Asymmetry in Responsible Investing Preferences and Beliefs

Jacquelyn Humphrey
Business School, University of Queensland
j.humphrey@business.uq.edu.au
+61 7 3346 8106

Shimon Kogan
Arison School of Business, Reichman University and The Wharton School
skogan@runi.ac.il
+972 9 9602741

Jacob Sagi
Kenan-Flagler Business School, University of North Carolina, Chapel Hill
jacob_sagi@kenan-flagler.unc.edu
919-962-8399

Laura Starks
McCombs School of Business, University of Texas, Austin
laura.starks@mccombs.utexas.edu
512-471-5899

February 2024

ABSTRACT: Empirical stylized facts in the literature concerning "sin" versus "angel" stocks display asymmetry. Through an experiment, we examine whether such biases can be micro-founded via individuals' preferences and belief formations. We find that negative environmental and social externalities have thrice the impact of positive externalities on investments, ceteris paribus. Further, negative externalities modestly increase pessimism about investment prospects while positive externalities have no discernible impact. The asymmetry is pervasive, heterogeneous, and comparable to the magnitude observed in loss-aversion. Beyond rationalizing stylized empirical facts, our findings help direct the growing theoretical literature that models the implications of non-pecuniary individual investor behavior.

NBER working paper 29288.

JEL Classification: G11, G41, C91

Keywords: responsible investing, sustainable investing, ESG, experiment

---

The authors thank Rob Bauer, Peter Bossaerts, Shaun Davies, Barry Oliver, Sébastien Pouget, Saphira Rekker, Paul Smeets, Gabby Walters, Paul Yoo, Josef Zechner and participants at the Inaugural Behavioral, Biological, and Experimental Economics conference, Penn State University, the Principles for Responsible Investment Academic Network conference and seminars at the Financial Research Network, University of Texas at Austin, Villanova University, Victoria University of Wellington, University of Geneva, the Wharton School and University of Woollongong for helpful comments. The authors are grateful to Terry Pan (Zheyao), Lee Seltzer, and Mitch Towner for research assistance; Larissa Garcia, Matt Pearsall, Jacob West and their team of assistants for lab assistance; and Alan McCrystal, Alex Cameron, and Paige Ottmar for programming assistance.

I. Introduction

On the face of it, the empirical literature on responsible investment (RI) features many examples where investors respond more sharply to negative firm externalities than positive ones. For example, although much of the RI industry has evolved to include positive ESG tilting, negative screening remains a pervasive strategy (e.g., Dimson et. al. 2020). A number of papers have shown that mutual fund investors react more to negative rather than positive ESG information (e.g., Krueger 2015; Bialkowski and Starks 2018; Hartzmark and Sussman, 2019). Moreover, evidence shows that sin stocks trade at a significant discount, but researchers disagree on whether a premium exists for angel stocks or green bonds (Hong and Kacperczyk 2009, Larcker and Watts, 2020, Bolton and Kacperczyk, 2021, Flammer 2021).

Despite the great interest in responsible investment in practice and in the academic literature, we are not aware of work demonstrating that investors' approaches to firms associated with positive versus negative social externalities are fundamentally asymmetric. A potential reason for this lack of analysis is that to demonstrate asymmetric behavior a researcher must control for the magnitude of negative versus positive externalities and establish that the reaction to one is disproportionate relative to the other. This is inherently challenging in the typical empirical setting. For instance, it is possible that, on average, the negative externalities exhibited by firms are larger, in absolute terms, than the magnitude of positive externalities. It is also possible that the risks associated with negative and positive externalities are asymmetric.

In this paper, we show experimentally that asymmetry is large, pervasive, and fundamental to non-pecuniary investor preferences. Our experimental framework isolates non-pecuniary preferences and allows us to vary the sign of the externality while keeping the magnitude of the stimulus constant. We exploit this design to establish that investor response to negative externalities is three times or more greater than the response to equivalent positive externalities. Our findings can therefore rationalize the empirical evidence of asymmetry found in many studies and across different domains, all through a single explanation.

In our experiment, which is adapted from Kuhnen (2015), subjects allocate money between a safe asset and a risky stock – the stock's expected return can either be high or low; i.e., it can have a 2/3 or a 1/3 probability of doubling in value, and 1/3 or 2/3 probability of halving. These parameters are fixed within a "trial" of six investment rounds and imply that the stock always has a weakly higher payoff than the safe asset, but uncertainty over payoffs remains. In each of the six decision rounds within a trial, subjects update their allocations to the stock and report their beliefs about the stock's potential payoffs.

The initial trials provide our baseline results for later comparisons in the experiment when we include responsible investment considerations. That is, after obtaining the baseline results, we ask the subject to identify a social cause important to them and we then link payments from the stock's payoff to that social cause. In the Positive (Negative) externality treatment, we increase (decrease) the donation to a preferred cause by the subject's gain from the stock allocation.

Notably, these donations do not affect the financial rewards received by the subjects. Because the donations relate solely to the stock payoff we can assess the variation in stock allocations across treatment conditions to gauge the influence of responsible investing preferences on the subjects' investment decisions. Subjects focused solely on personal gains, i.e., their own payoffs, should allocate an identical amount across all treatments. Alternatively, those who care about the externalities resulting from their investments would be expected to consider how allocating to the risky asset affects both their individual profits and the social cause they have chosen to support.

This design affords two important advantages. First, we are able to keep the magnitude of the externalities constant across treatments, which cannot be done in the field. Thus, we can compare behavior not just against the base case treatment, i.e., without externalities, but also between the Positive and Negative externality treatments in a way that quantitatively controls for the external stimulus. Second, the treatments only distort investors' non-pecuniary motives and, in so doing, allow us to study how individuals may incorporate these metrics into their investment decisions. Under neoclassical assumptions, allocations and beliefs should not be affected by the treatments. That is, the design rules out common alternative financial motives for RI preferences, such as expectations about higher returns or lower risk.

The experiment generates multiple rounds of data on individuals' allocations and beliefs. We document a statistically strong effect from the negative RI externalities – reported beliefs and allocations are significantly different from the baseline. Moreover, the influence on allocation decisions is economically pronounced (roughly a 30% reduction in average allocation across subjects). However, we find a much smaller effect from the negative externality treatment on the individuals' beliefs formation – roughly a 2% reduction in the perceived probability that the stock's returns are positive. In stark contrast, for positive RI externalities we find far less internalization on either allocation decisions or beliefs. The asymmetric impact of negative externalities on preferences and beliefs in our experiment is similar in magnitude to asymmetries observed in studies of behavior in loss versus gain frames (e.g., loss aversion).[1]

One potential concern with our experimental design could be the validity of the mapping between the Negative treatment to choices in the field. In the Negative treatment, payoffs from investments are associated with a reduction in the contribution to a nonprofit that works towards positive externalities. Thus, one could question whether this design actually captures the choice investors face in practice. In particular, sin stock investments can result in increased

[1] Somewhat puzzling are the weak Positive treatment effects we observe because a simple charitable-giving motive has been prevalent in both practice and in experimental economics, which suggests that an investment context may mask or muddle charitable-giving behavior even while clearly showing that the negative RI stimulus results in a pronounced aversive response.

payoffs from harmful activities rather than "less beneficial" activities. To address this potential concern, we design a subsequent survey that uses the same set of causes as in the main experiment and ask subjects to state their preferences between contributing funds towards a negative cause versus reducing the funds of a charity opposing that cause. Respondents of similar demographics to the original experiment's population exhibit a strong preference for reducing funding from a charity doing good rather than contributing funds that will be used to do harm. Thus, the results in our main experiment likely *understate* the degree of asymmetry in the field.

Beyond the average treatment effect, we find significant heterogeneity across subjects' RI sensitivities to negative but less so to positive externalities. Roughly one third of subjects exhibit a pronounced bias against the Negative treatment relative to their allocation response in the Positive treatment. The remaining two thirds may be said to increase their allocation to the stock in the Positive treatment by roughly the same amount as they decrease their allocation in the Negative treatment, i.e., their treatment effect is symmetric. Thus the asymmetry we find appears to be driven by a subset of the population, albeit it a substantial one. In an Appendix, we use within-subject data and a reduced-form structural model to estimate heterogeneous allocation curves as functions of investors' reported probabilities in the three treatments. These estimates can potentially be used to construct or calibrate more descriptive theoretical models RI of behavior.[2]

Our paper builds on, and contributes to, the extensive literature that demonstrates the prevalence of asymmetry in individual decision making across multiple domains. The most prominent example is the Kahneman and Tversky (1979) concept of loss aversion, in which individuals are willing to take significant risk to avoid perceived losses, while concurrently

[2]Most models of non-pecuniary RI investing are limited to two types of investors: Those with only financial concerns and those that additionally exhibit non-pecuniary preferences. The latter are typically archetypes who either screen stocks (extreme asymmetry) or exhibit no asymmetric RI preferences.

adopting a more conservative approach when their choices are framed as gains.[3] Correspondingly, a complementary body of literature documents asymmetry in individuals' information processing. This literature also provides examples of preferences influencing individuals' belief formation.[4] More pertinent to our own contribution and mirroring preference asymmetry in loss-aversion, Kuhnen (2015) demonstrates that individuals tend to over-pessimistically update their beliefs when payoffs are framed as losses rather than gains.

Although evidence exists regarding individuals' RI choice behaviors and belief formations based on their non-pecuniary preferences, direct evidence addressing asymmetry is lacking.[5] This gap raises the question of whether asymmetry in individual RI preferences can serve as a basis for observed aggregate asymmetries. That is, the absence of a direct link, and considering the rich history of asymmetry in individual behavior found in other domains, provides context and motivation for our work. Exploring the role of asymmetry in RI preferences is important for at least two reasons. First, it can offer insights into reconciling the disparate findings related to aggregate asymmetries in RI. Second, it can help in distinguishing how actual investor behavior aligns with existing and future theoretical models of RI preferences.[6]

The paper is structured as follows. In Section II we outline the experiment, reporting the details of the analyses and their interpretations in Section III. In Section IV we discuss our

[3] Other examples include the disparity between the valuation of an owned object and the price at which one is willing to acquire it (Kahneman, Knetsch, and Thaler, 1990); the more pronounced vividness of negative versus positive memories (Kensinger and Schachter, 2008); and the greater influence of anticipated regret/disappointment versus rejoicing/elation in emotion-based choice (Mellers, Schwartz, and Ritov, 1999).
[4] For example, in the well-known "confirmation bias" individuals place more weight on new information that confirms their beliefs – especially when those beliefs are favorable to their worldview. See, also, Section 9 in the review article by Benjamin (2019).
[5] For evidence, see, for example, Riedl and Smeets (2017), Hartzmark and Sussman (2019), Barber, Morse and Yasuda (2021), Bauer, Ruof, and Smeets (2021), Geczy, Jeffers, Musto and Tucker (2021), and Yoo (2022).
[6] Among the theoretical RI models, Heinkel, Kraus and Zechner (2001), Luo and Balvers (2017), Pedersen, Fitzgibbons and Pomorski (2021) and Zerbib (2022) allow for asymmetric RI preferences (e.g., negative screening). Other prominent studies do not (e.g., Chowdhry, Davies and Waters 2019; Oehmke and Opp 2023; Pastor, Stambaugh and Taylor 2021; Goldstein, Kopytov, Shen, and Xiang 2023). None consider the influence of RI preferences on beliefs.

paper's contribution to the literature in light of prior work. In Section V we provide our conclusions.

II. Experimental Design

A. Description of Experiment

The experiment is organized around a basic set of tasks we term a "trial" performed through a computer terminal.[7] The experiment itself consists of a series of these trials, some of which include the positive or negative externality treatment effects. Before describing the experiment's sequence of events, we detail the mechanics of a single trial.

At the beginning of each trial, participants are informed that a stock in which they can invest may be one of two types: a high payoff stock that doubles the amount invested with a probability of 2/3 and halves the investment with a probability of 1/3; or a low payoff stock that doubles investment with a probability of 1/3 and halves it with a probability of 2/3. Thus, the expected return on the stock ranges from 0% to 50%, based on whether one believes they are facing the high or low payoff stock. Participants are also informed that at the beginning of the trial the computer randomly selects, with equal probability, the stock type and that the stock's type remains fixed for the duration of the trial. The trial consists of six rounds of investment. At the end of each round, the randomly realized return of the stock is shown, providing information about the stock type.

Before the first round starts, participants are asked to estimate the probability that they are facing the high payoff stock. The correct answer is 50% as participants were told that the computer randomly selects the stock type with equal probability. Participants next allocate an endowment of 100 experimental currency units (ECU) between the stock and cash.[8] A snapshot

[7] Each subject participates in the experiment through a separate computer terminal, asynchronously and independently of other subjects. We use Qualtrics. Randomization by the software is independent across subjects.
[8] 100 ECU is equivalent to US$10.

of the initial round allocation screen is displayed in Panel I of Appendix A. The computer then randomly generates an outcome consistent with the distribution of the stock given its true type, and participants receive a report of the results of their investment round, i.e., whether the stock doubled or halved, as well as the value of their winnings. Having observed whether the stock doubled or halved in the first round, participants are asked to again estimate the probability that the stock is of the high payoff type and to allocate 100 ECU between the stock and cash. Thus, the only thing that changes from one round to the next within a trial is the information that the subject has about the stock and, potentially, the subject's allocation to the stock. This process repeats until six investment rounds are completed. Panel II of Appendix A depicts a screen capture of what a subject might see after round three of a trial.

To encourage attentiveness, the experiment incorporates prompts asking participants if they are sure of their decisions whenever they appear to violate a monotonicity condition. For example, a prompt appears any time a stock outcome "halves", but a participant *increases* either the allocation to the stock or the estimated probability that it is a high payoff stock.

We now describe the full sequence of events (and trials). At the start of the experiment, participants are told that they will be taking part in an experiment in decision-making, and their main task will be to choose how to allocate an investment of 100 ECU between a risky stock and cash in each of a set of rounds. Participants are told that the total payout they can expect from the experiment consists of a US$7 participation fee, plus the total stock and cash payoff from one randomly selected non-practice round, plus US$1 if the stock-type probability estimate made by the participant in the randomly selected round is within 5% of the true probability. Participants are also told that, given the stock's history, there is an objectively true probability that the stock is the high payoff type.

Each subject participates in seven trials divided into four blocks. The first block consists of a single practice trial and serves to familiarize participants with the process. The second block consists of two trials that set a baseline we term the Neutral treatment. A subject then participates in two additional treated blocks, each consisting of two trials: a Positive treatment block and a Negative treatment block, in randomly determined order. Because each trial consists of six investment decisions and probability elicitations, excluding the practice trial, we collect 36 observations of allocation decisions and 36 observations of likelihood perceptions, per subject.

Treatment proceeds as follows. Once participants complete the Neutral block, we elicit their social preferences by asking them to rank six social issues in order of importance. The six issues are: animal welfare, environment, refugees, poverty, human trafficking, and gender discrimination. Participants then view a screen that describes two nonprofit organizations related to their top ranked social issue. They are asked to select one of the two nonprofit organizations to link to their trading outcomes. This process is repeated for the second-ranked social issue. Panel III of Appendix A depicts a snapshot of the social issue decision screens.

Participants subsequently proceed to either the Positive or Negative treatment block. In the Positive block, participants are told that an amount of money equal to their stock payoff *would be donated* to their selected nonprofit by the experimenters. In the Negative block, they are told that an amount of money equal to their stock payoff *would be deducted* from their selected nonprofit's donation account.[9,10] It is important to note, as emphasized to the participants, that the amounts donated to (or deducted from) a selected nonprofit would not affect the participant's own gains during the experiment. The nonprofit remains fixed for both

[9] We randomize whether the first or second ranked social issue's non-profit is assigned to the Positive or Negative block.

[10] For each subject, the randomly drawn payment round may be from the positive or negative condition and associated with one of the subjects' two chosen nonprofits. We therefore aggregate the specified donations from each of the participants' outcomes across the non-profits, netting the positive and negative payments. For example, Rainforest Alliance was the non-profit associated with the payment round for 12 participants. The payment amounts were -$3.90 and $49.00 from the negative and positive blocks, respectively. Therefore, the final payment made to the Rainforest Alliance by the experimenters was $45.10. The net payment was negative for four of the twelve non-profits, and these payments were taken to be zero. Receipts for all payments were uploaded to a Dropbox folder and subjects were subsequently emailed a link to this folder.

trials of a given block, but changes across treatment blocks. Panels IV and V of Appendix A show the instructions for the Positive and Negative blocks, respectively. Also shown in Panels IV and V of Appendix A is that during each trial, participants receive a report of the amounts to be potentially gained by, or deducted from, the nonprofit in past investment rounds (this can be compared with the feedback provided in the Neutral block – see Panel B).

The six non-practice trials in the experiment are payoff-equivalent for the subjects regardless of treatment. The only difference across blocks is that investment decisions may determine whether a nonprofit to which the subject exhibits some affinity gains (losses) money in the Positive (Negative) treatments. This allows us to examine the causal impact of treatment on likelihood perceptions and willingness to invest (i.e., preferences) given a likelihood perception.

A. Description of Subjects

We recruited 160 participants from the University of Texas at Austin (62 identified as male, 97 as female, one did not identify themselves). These participants received an average of $18.30 for their participation and choices (range of $12-$28, including a show-up fee of $7). The age of the participants ranged from 18 to 34 with a median of 20. This age group merits particular attention because of prevailing interest in the potential effects of millennials on asset markets.[11] During our study period the definition of a millennial is an individual whose age varies from 23 to mid-30s. Most participants were students at the school, with 50 enrolled in business-related degrees, 39 in natural sciences, 19 in medicine, 16 in engineering, 10 in social sciences, and the remainder in arts/humanities, law, nursing, mathematics and communication.

[11] In particular, the increasing availability of RI products are often viewed as a response to anticipated demand from millennials, who are expected to receive large transfers of wealth from baby boomers in coming years.See, for example, https://pewrsr.ch/2Op4i3b; https://go.ey.com/2XvjCiP; and https://bit.ly/2O1r5mS.

Somewhat surprisingly, demographic variables were not associated with the main effects we identify in the following section, and we omit them from our reported analyses.

One way of gauging subjects' comprehension and cognitive abilities around probability estimates is to ask whether their reported probability is 50% before observing any signals. Indeed, in 88% of the trials the subjects report the correct first-round probability, with this fraction being stable across treatments. Alternatively, 86% of subjects responded correctly on the first round always or almost always.

III. Empirical Results

We focus on two variables, the participants' preferences for investment allocations and their beliefs, that is, the participants' allocations to the stock investment (in ECU) and their reported probability estimates of whether the stock's type is high payoff. To analyze the treatment effect on each of these variables we begin with a simple comparison of their average levels across the three treatments (Negative, Neutral, and Positive). This is a very conservative use of the data as it produces a single observation for each subject-treatment, thus treating responses within it as being perfectly correlated. Further, it lacks any of the controls that are later useful in isolating the treatment effects.

In Table 1 we report the participants' allocations in Panel A, showing the differences across treatment conditions in the allocations to the risky stock. On average, subjects allocated 28.1, 36.7, and 39.0 (all out of 100 ECU) in the Negative, Neutral, and Positive treatments, respectively. Thus, relative to the Neutral treatment, the average allocation to the stock is 23% lower in the Negative treatment, but only 6% higher in the Positive treatment. Moreover, the difference between the Neutral and Negative conditions is statistically significant at the 1% level, while the difference between the Neutral and Positive conditions is only marginally significant at the 10% level (Panel B).

Table 1: Stock Investment Allocations and Probability Estimates

This table reports the means, standard deviations, and ranges of the participants' choices regarding the stock investment allocations in each of the conditions and their estimates of the probability that they were facing a high payoff stock. The table also includes a test for whether the allocations and estimates are different across treatments, and a test for the effect asymmetry by calculating the ratio of allocation between the neutral minus negative treatments and positive minus neutral treatments ("asymmetry multiplier"). All *t*-tests are based on matched samples at the subject-treatment level; *p*-values are reported in Panels B, C and E.

Panel A: Stock Allocation in ECU

	Negative (N=160)	Neutral (N=160)	Positive (N=160)	Total (N=480)
Mean (SD)	28.104 (17.817)	36.651 (20.557)	39.044 (21.739)	34.600 (20.606)
Range	0.000 - 100.000	4.167 - 95.833	2.833 - 100.000	0.000 - 100.000

Panel B: *p*-values from Matched Sample *t*-tests on Stock Allocation

	Negative=Neutral	Negative =Positive	Positive = Neutral
Probability	0.0%	0.0%	7.1%

Panel C: Asymmetry in Stock Allocation ECU

	Neutral - Negative (N=160)	Positive - Neutral (N=160)	Asymmetry Multiplier	Test: Multiplier = 1 (prob)
Mean (SD)	8.547 (18.570)	2.392 (16.630)	3.5727	0.0081

Panel D: Probability Estimates

	Negative (N=160)	Neutral (N=160)	Positive (N=160)	Total (N=480)
Mean (SD)	46.713 (13.039)	48.976 (11.492)	48.618 (12.327)	48.102 (12.316)
Range	9.667 - 75.917	5.833 - 77.417	8.833 - 81.667	5.833 - 81.667

Panel E: *p*-values from Matched Sample *t*-tests on Probability Estimates

	Negative=Neutral	Negative =Positive	Positive = Neutral
Probability	8.8%	12.3%	76.4%

In Panel C we calculate the asymmetry of the treatment effects. The magnitude of the treatment deviations from the baseline allocation is more than 3.5 times larger for "Neutral minus Negative" than for "Positive minus Neutral" and this asymmetry multiplier is significantly different from one at the 1% level. Since we have designed the study so that the magnitudes of the positive and negative stimuli are identical, the profound difference in subjects' reactions reflects a strong asymmetry in the valence with which these stimuli are perceived. For comparison, it is worth noting that experimental studies of loss aversion document an asymmetry multiplier that varies between 1.43 and 4.8 (Abdellaoui et al., 2007).

Moreover, if we separate our results for causes that the subject ranked first versus second (not shown in the table), we find differences that are consistent with our effect being causally driven by considerations of externalities. That is, we find that allocations are more affected by the treatment when it is randomly linked with a cause that the subject ranked first, compared to when it ranked second (27.7 vs. 30.5 ECU in the Negative treatment, 41.0 vs. 37.9 in the Positive treatment).

In Panel D for each of the treatments we report the average level of the participants' estimates of the probability that the stock is the high payoff type. In comparing Panels D and A, it is evident that the treatment effects are not as strong for probability estimates as for allocations to the risky stock. Probability estimates across the three treatments are similar, being 46.7%, 49.0%, and 48.6% for the Negative, Neutral, and Positive treatments, respectively. Testing for the differences across these probability estimates, which we report in Panel E, we find that only the Negative and the Neutral treatments are (marginally) significantly different with no significant difference between the probability estimates for the Positive and the Neutral treatments. It should also be noted that the magnitude of probability differences between Negative and Positive (or Neutral) treatments are roughly consistent with the differences found in Kuhnen (2015) across probability estimates elicited in treatments of loss versus gain frames.

Table 1 provides a first glimpse of results that turn out to be robust in the experimental data: Subjects' allocation choices and probability estimates are far more sensitive to the Negative than the Positive treatment. That is, subjects in the Negative treatment sacrifice their own financial gain in the experiment to keep from harming the charities they selected. Moreover, the sensitivity to the Negative treatment is more pronounced for allocation choices than for probability estimates. It seems quite surprising that the response to Positive externalities is relatively muted. Put differently, subjects in the Positive treatment do not appear especially eager to raise the payoff of their chosen nonprofit by taking on greater risk, relative to the Neutral treatment. This result is puzzling because charitable giving is both prevalent in practice and in experimental economics (e.g., see the meta-analysis of the Dictator Game in Engel, 2011). Consider that, in our design, *any* increase in allocation to the risky asset in the Positive treatment makes the associated nonprofit strictly better off.[12] Conventional wisdom, and intuition, suggests` that a significant number of subjects would therefore allocate more to the stock in the Positive treatment than they might otherwise. The fact that they do not signals that subjects respond to the investment context in a manner that is different from conventional charitable giving.[13] We interpret this as evidence that the presence of risk masks the charity giving context and creates a different frame in the minds of subjects. The finding that subjects are sensitive to the imposition of negative externalities is remarkable precisely because, in the same context, they are nearly indifferent to charitable giving.

A. Allocations

Although the allocations are affected by subjective probabilities, the summary statistics in Table 1 give evidence that RI considerations affect allocations to the stock above and beyond what can be easily explained by shifts in beliefs alone. To help disentangle the effects, we focus on the allocation decision given subjects' likelihood perceptions (i.e., preferences), and later separately examine the effects of RI on probability estimates (i.e., beliefs).

Before applying a more structured approach to the data, we filter to retain only subjects with at least weakly rational behavior. We require that a subject in the Neutral treatment should,

[12] Even if subjects may not recognize this fact in the abstract, they can perceive it in the report we provide (e.g., Panel D, Appendix I) on how their allocations in the Positive treatment translate into outcomes for the nonprofit.
[13] List (2007), for instance, provides examples of contexts that can greatly influence pure charitable-giving motives via the Dictator Game.

on average, exhibit an at least weakly positive relationship between actual and subjective (i.e., estimated) probabilities that the stock is high payoff, and an at least weakly positive relationship between subjective probabilities and stock allocations. We interpret violations of these conditions in the Neutral treatment to signify lack of engagement or confusion about the basic experimental tasks. We test for the conditions, at the subject level, using linear regression. Of the 160 subjects, 35 were dropped because they violated one or both of these requirements. It is important to note that this filter is only applied through the Neutral treatment and thus does not implicitly condition on behavior in the main treatment cells of interest – the Positive and Negative conditions. We proceed by analyzing the data from the remaining 125 subjects though it bears emphasizing that we obtain qualitatively similar results without this filter.[14] We term this group of subjects the rational subsample.

To control for the impact on probability assessment and examine the treatment effects on allocation separately from their effects on beliefs, we pool allocation observations across all rounds based on subjects' *reported* (i.e., subjective) probability bins. These are depicted in Figure 1. A number of suggestive patterns emerge from the plot. First, allocations in the Negative treatment are lower across almost the entire range of subjective probabilities compared with the Neutral or Positive treatments. Second, the effect does not appear to be uniform – instead, the difference between the allocations in the Negative treatment and the other treatments appears to increase as subjective probabilities increase. This is especially striking when plotting differences from the Neutral treatment (Panel B) where one sees both the asymmetry in response to the treatments as well as the subjective probability dependence.

[14] Together with the subject prompts (in the event of monotonicity violations), the filter effectively acts as a control for comprehension. The fact that the results are largely robust to explicitly controlling for these effects is reassuring.

Finally, there appears to be only a marginal treatment effect on allocations when comparing the Positive condition to the Neutral one.[15]

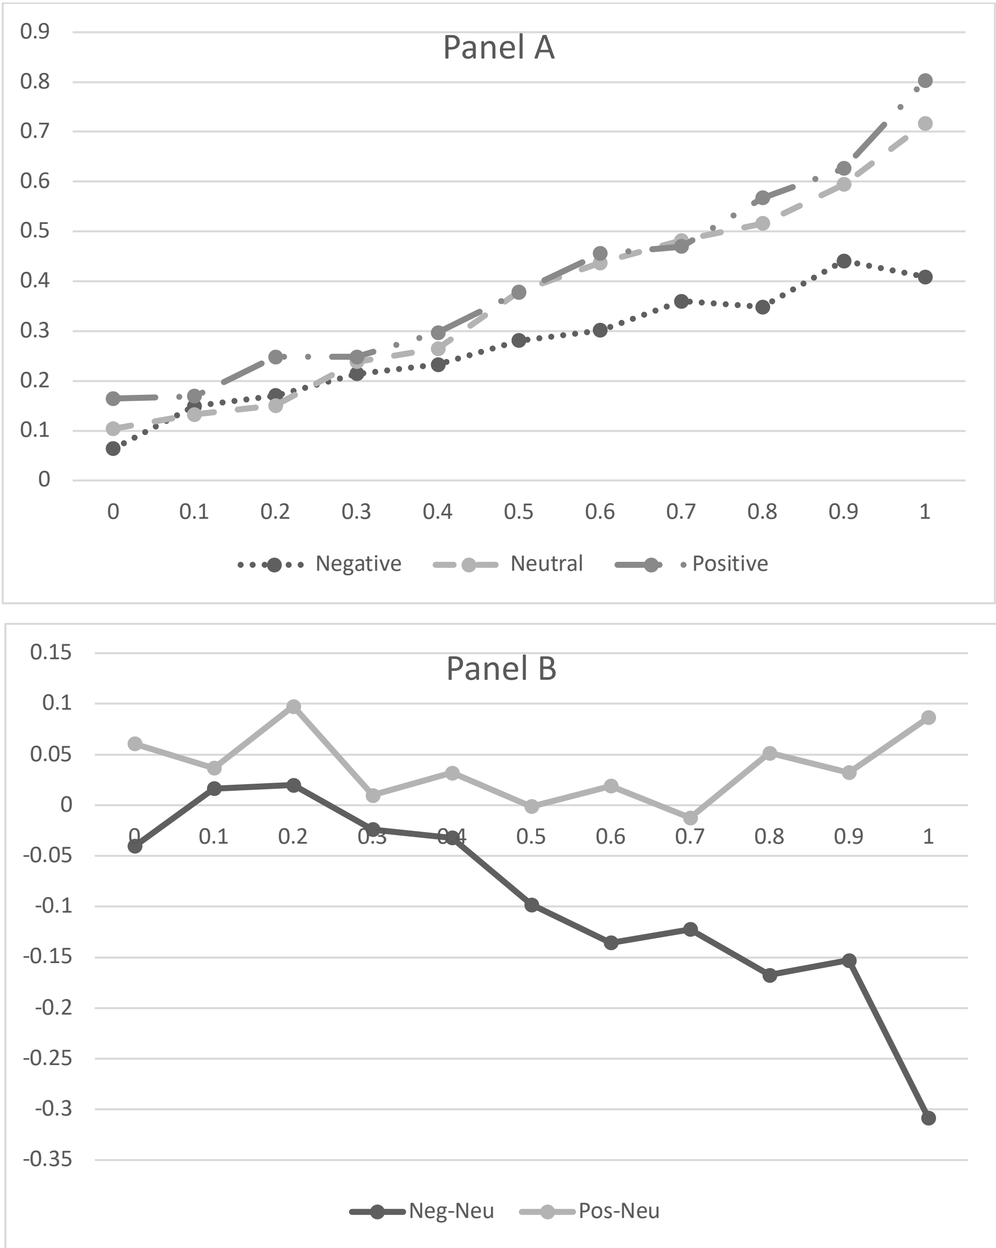

Figure 1: Treatment Effects on Allocations

This figure shows the treatment effects on allocations (y-axis), scaled to 0.0-1.0, against subjective probability estimates that the stock is of high type (x-axis). We group observations based on the subjects' reported probabilities rather than the objective probabilities to control for allocation differences that may arise from different assessments of probabilities. Panel A plots the average allocation across the three treatments while Panel B plots the average difference in allocation in the positive and negative treatments, both relative to the neutral treatment.

[15] The average Neutral treatment allocation per bin is greater than that of the Positive treatment in only two of eleven plotted bins. Under a 50-50 null (consistent with indifference) between the bin statistics, this is associated with a two-sided *p*-value of 6.6%.

To further quantify these observations, we regress allocations in each round on treatment dummies, the reported subjective probabilities, and interactions between treatment dummies and the reported probabilities. All regressions include subject fixed effects to control for heterogeneity in average allocations across subjects. In Table 2 we report the regressions of each round of the subjects' percentage stock allocations on their reported subjective probabilities, treatment dummies, and interactions between the reported probabilities and the treatment dummies. The results confirm the patterns reported in Table 1 and Figure 1. Without the probabilities included, Column (1) shows that allocations in the Negative treatment are, on average, a highly significant 9.6 ECUs lower than in the Neutral treatment. In contrast, the increase in allocation observed in the Positive treatment is 1.6 ECUs higher than in the Neutral treatment. When we add the probabilities, in Column (2) the results demonstrate that, while there is a strong treatment effect, it does not appear to have a highly significant constant component (the un-interacted treatment dummies). On the other hand, we find that the response of allocations to probabilities is much lower in the Negative treatment relative to the Neutral treatment, 0.381 (=0.589-0.208) vs. 0.589, but there is no significant difference between the Neutral and Positive treatments. This is consistent with a more pronounced allocation reduction in the Negative treatment, relative to the others, as subjective probabilities increase.

Because we have multiple observations per participant, the greater number of observations in the analysis reported in Table 2 (as compared to Table 1) permits a more precise estimate of the asymmetry multiplier (the ratio of the magnitudes of the deviations for the Negative and Positive treatments from the Neutral treatment). Using estimates from the regression parameters, the multiplier in Column (1) is 6 (=.096/.016) – corresponding to substantial statistical and economic significance.

Table 2: Stock Allocations and Probabilities

The table reports regressions of round-by-round percentage allocations to the stock on reported subjective probabilities ("Prob"), treatment dummies, and interactions between reported probabilities and the treatment dummies. The Neutral treatment is the baseline. Standard errors are in parentheses. All regressions include subject fixed effects. The analysis is limited to subjects exhibiting at least weak rationality in the Neutral treatment. The asymmetry multiplier is calculated as the ratio between the allocation in the Neutral-Negative treatment and the Positive-Neutral treatment (in column 2, the ratio is evaluated at the median probability of 50%). Asterisks denote the probabilities for statistical tests of differences from zero (*p<0.1; **p<0.05; ***p<0.01).

	(1)	(2)
Negative Block Dummy	-0.096***	0.012
	(0.008)	(0.017)
Positive Block Dummy	0.016**	0.030*
	(0.008)	(0.017)
Prob		0.589***
		(0.023)
Prob x Negative Block Dummy		-0.208***
		(0.032)
Prob x Positive Block Dummy		-0.026
		(0.033)
Observations	4,500	4,500
R^2	0.369	0.527
Asymmetry Multiplier	6.00	5.41
Test of asymmetry multiplier	0.000	0.000

The results from Column (2) in Table 2 may be interpreted as evidence that the channel through which the Negative treatment primarily impacts subjects' preferences is their sensitivity of allocation to subjective probability: The greater the probability of a high stock payoff, the greater the differences in the allocations between the Negative treatment and the other treatments. Note that a significant "charitable-giving" motive implies a significant Positive block dummy coefficient. This is because any additional amount allocated to the stock in the Positive treatment results in some distribution to the nonprofit (even when the subjective probability is zero that the stock is high type). That said, consistent with the absence of a charitable-giving motive, allocations in the Positive and Neutral treatments are also similar in their sensitivity to the subjective probability of a high payoff. Note that adding the subjective probabilities to the regression does not change the multiplier very much as the multiplier calculated in Column (2) of the table, which is evaluated using the median subjective probability of 50%, is similar in magnitude to that reported in Column (1).

B. External Validity and Survey Evidence

The Positive treatment in our experiment mimics certain RI objectives in that it creates "good" externalities. Although the Negative treatment does not directly create harm, it can be viewed as creating less good. However, this may raise questions as to whether our experiment measures an asymmetry that is meaningful in practice. Given the institutional restrictions that rule out a design that could be plausibly linked to actual harm creation, we conduct a separate survey to test whether our results *understate* the strength of asymmetries. That is, we hypothesize that subjects' negative reactions to funding a group that advocates against environmental or social causes would be stronger than their negative reactions to taking funds away from a group that advocates for these causes. In this subsection, we report on the results of a separate survey that strongly support this hypothesis.

The survey was run on Prolific, pooling from the general US adult population. To match subjects' age and gender demographics in the main experiment, we specified that participants be between the ages of 18-25 and requested an equal proportion of males and females. In all, 108 participants responded to the survey.

To mirror the design of the main experiment, at the start of the survey, participants are asked to rank the six available causes, which were identical to the ones used in the main experiment. For their highest ranked cause, each subject is asked to make a hypothetical choice between two negative effects: (i) providing funds to a firm that lobbies against their chosen cause (i.e., harm creation) and (ii) cutting funds from a charity that lobbies for their chosen cause. Subjects are also given the option to indicate that both choices are equally bad. An example of this choice for a social cause concerned with poverty is provided in Appendix B.

Table 3 summarizes the main survey results. Among the 108 responders, 36% preferred the choice of cutting funding to the initiative that would promote their cause and only 11% preferred the choice of contributing to the initiative that would harm their cause. The remaining 53% were indifferent, that is, they indicated that they thought both choices were equally bad. Using a simple Chi-squared test, at a 1% level we can reject the null that the survey participants' responses are uniformly distributed across the three options provided. This allows us to reject the null hypothesis that most respondents randomized their selections. Likewise, we can reject the null that respondents were equally disposed towards reducing good versus creating harm: Three times more subjects preferred to reduce good than create harm. Viewed another way, out of the respondents who indicated a preference, 76% selected the reducing good choice over the creating harm choice.

At the end of the survey, participants were presented with a few exit questions, including a self-evaluation of their level of engagement (on a scale of 1-10). This provides an indirect proxy for their level of interest in these RI issues overall. We divide the subject pool into those that assessed their engagement as "10" (61%) and the remainder (39%). While the maximally engaged group exhibits a far greater propensity to reduce good than create harm, remaining subjects still exhibited a significant preference in the same direction.[16]

The results of the survey strongly support the thesis that the subjects in the main experiment would have exhibited even greater aversion to the negative treatment if it had incorporated harm creation in the design. In other words, our estimates from the main experiment likely *understate* the true asymmetry in RI preferences.

[16] We also included a task comprehension question at the end of the survey. Of the 108 participants, all but 2 answered the comprehension question correctly. Moreover, 87%/92% of subjects rated the survey as being very engaging/clear (giving it a score of 8 or more out of 10).

Table 3: External Validity Survey

The table reports survey results testing whether subjects would choose to "reduce good" by cutting funding from a charity lobbying to ameliorate a social issue they rank as important, "create harm" by donating to a firm lobbying to exacerbate the social issue or are indifferent between cutting and donating. The first Chi2 column reports the *p*-value for a Chi-squared test for uniformity of responses (i.e., randomization). The second Chi2 column reports the *p*-value for a test of the null that the propensity to cut is the same as to donate in the population. The subsample of subjects who are "Maximally engaged" assessed their engagement in the survey as "10 out of 10".

	Cut	Donate	Indifferent	Total	Chi2(Uniform)	Chi2(Cut=Donate)
Total	39	12	57	108	0	0
Maximally engaged	24	5	37	66	0	0
Others	15	7	20	42	0.0464	0.088

C. Heterogeneity of Treatment Effects

The results so far focus on averaging treatment effects across subjects. An interpretation of the results in Table 2 in terms of an impact on individual preferences is, at best, indirect because the analysis relies on the average treatment effect across all subjects. In other words, Table 2 reports aggregated choices but is not directly informative about the behavior of the average individual. The primary focus of this subsection is to quantify the heterogeneity of RI preferences in the subject pool. We investigate whether the average treatment effect is broad-based or attributable to a few subjects with very strong social preferences. Arguably, a prevalent effect can be more compellingly extrapolated to the general population. In addition, the prevalence and strength of the effects can be potentially useful in calibrating models that both incorporate social preferences and include more than one investor type.

We begin by considering the asymmetry of the treatment effect within subjects. For each of the 125 rational subjects, we want to compare the Positive treatment effect to the Negative treatment effect while holding constant the probability that the stock is "high payoff". For each subject, i, and each objective probability realization, q, we calculate the *asymmetry score* as $2\overline{ECU}_{i,q,Neu} - \left(\overline{ECU}_{i,q,Neg} + \overline{ECU}_{i,q,Pos}\right)$. Here, $\overline{ECU}_{i,q,Neu}$ is the average Neutral treatment stock allocation by subject i when q is the objective probability that the stock is a high type. We similarly define $\overline{ECU}_{i,q,Neg}$ and $\overline{ECU}_{i,q,Pos}$. Typically, $\overline{ECU}_{i,q,Neg} \leq$

$\overline{ECU}_{i,q,Neu} \leq \overline{ECU}_{i,q,Pos}$, so that one can think of the asymmetry score as a concavity measure. When it is greater than zero, the subject exhibits a stronger Negative than Positive treatment effect.[17] Of the 125 rational subjects, 38 exhibit a strictly greater than zero asymmetry score across *all* realizations of objective probabilities (i.e., regardless of the outcome-relevant information made available to them through observing the stock's history). We refer to these subjects as having a strong bias against investing in the risky stock in the Negative treatment. By comparison, not a single subject exhibited an asymmetry score greater than zero across all realizations of objective probabilities.

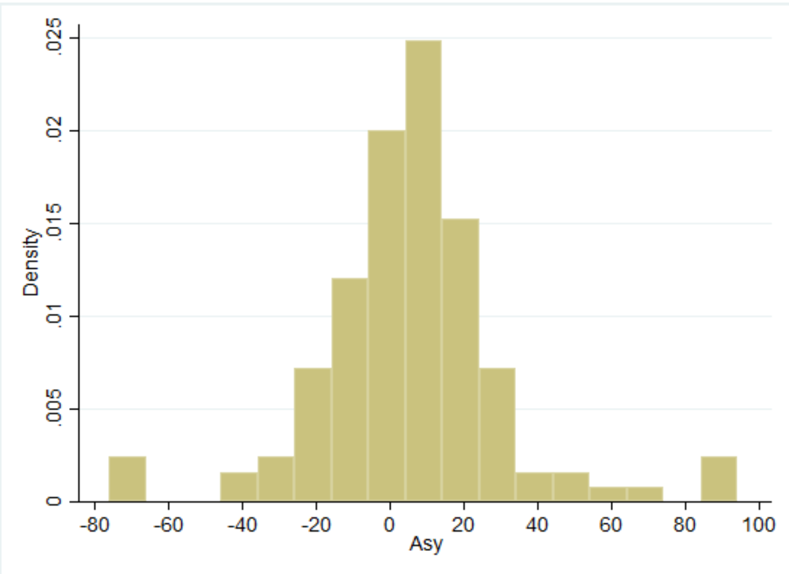

Figure 2: Asymmetry in subject reaction to the Negative and Positive Treatments
An asymmetry score that is greater than zero corresponds to a larger stock allocation deficit in the Negative treatment, relative to the Neutral treatment, than the excess stock allocation in the Positive treatment.

[17] We note that, at a given level of q, the asymmetry score can only be calculated for a subject that faces at least one realization of q in each of the treatments. This is guaranteed for q=½ because in the first round of every trial q=½, but the chance that an asymmetry score can be calculated diminishes for values of q further away from ½. For instance, no subject experienced five consecutive rounds of the high outcome in *each* of the treatments. So an asymmetry score for $q \approx 0.97$ or higher could not be calculated for any subject.

Next, for each subject, we average the asymmetry score across the objective probability realizations. Figure 2 plots the density of subjects' averaged asymmetry scores. The average and median asymmetry is 5.75 and 5.0 ECUs, respectively, which is roughly 17% of the corresponding Neutral treatment stock allocation. Eighty out of 125 subjects exhibit an averaged asymmetry score that is greater than zero, which indicates that there are roughly twice as many subjects with asymmetry scores greater than zero than those whose asymmetry scores are less than zero.

Table 4 reports the distributional properties of the asymmetry scores across the 125 rational subjects. In addition, the table reports the same statistics for two subsets of this sample: The 38 subjects exhibiting a strong negative bias against the stock in the Negative treatment exhibit a large and highly significant average asymmetry score of 25.84 ECUs – amounting to nearly 60% of their allocation in the Neutral treatment. The remaining subjects, by comparison, exhibit a statistically insignificant average asymmetry score. This crudely suggests that the observed sample-wide asymmetry across treatments is largely driven by roughly a third of the subjects who severely curtail their allocation to the stock in the Negative treatment.

Table 4: Distribution of Average Asymmetry Scores
The table reports statistics for the average asymmetry score of various subject subsamples: The full sample of 125 at least weakly rational subjects, the subsample of 38 "strongly biased" subjects who exhibit strictly greater than zero asymmetry regardless of their outcome-relevant information set, and the remaining 87 subjects. When the asymmetry score is greater than zero, the subject exhibits a stronger Negative than Positive treatment effect. Asterisks denote the probabilities for statistical tests of differences from zero (*p<0.1; **p<0.05; ***p<0.01).

Asymmetry score	N	mean	sd	p90	p75	p50	p25	p10
Full sample	125	5.75**	25.38	30	15.56	5	-6.73	-19
Strongly biased	38	25.84***	23.21	63.65	30	18.65	10	5
Not strongly biased	87	-3.02	21	15.56	9.82	-0.63	-12.36	-22.5

The above observations indicate that the subject pool can be plausibly modeled as being composed of two types. The first type increases their allocation to the stock in the Positive treatment by roughly the same amount as they decrease their allocation in the Negative treatment, i.e., their treatment effect is symmetric. The second type exhibits a pronounced bias against the Negative treatment relative to their allocation response in the Positive treatment.

While the analysis of heterogeneity presented above has the virtue of being relatively simple, it is subject to several potential shortcomings. First, in order to calculate the asymmetry measure, we resort to using the objective rather than the reported probabilities. Moreover, we effectively assume that asymmetry at each probability level is independent and contributes equally to asymmetry, within subject. Additionally, the analysis does not account for within subject experimental noise that may bias allocation statistics because the range of possible allocations is bounded. In Appendix C we address these concerns by estimating a subject-level stimulus-response model that incorporates experimental noise and all *reported*, rather than objective, probability observations. This more elaborate approach confirms the findings described above: A substantial proportion of subjects exhibit asymmetric treatment effects, with a large Negative treatment effect that increases with the probability that the stock is high type (see Figure C-1). In particular, after controlling for the various confounds addressed in Appendix C, the overall proportion of subjects that exhibits strong asymmetry appears to be closer to one half (see Table C-3).

Theory and intuition suggest that the magnitude and prevalence of the effects we find across subjects, if representative of the population, will impact asset prices (e.g., Heinkel, Kraus and Zechner, 2001; Zerbib, 2022). The relatively weak Positive treatment results and strong Negative treatment results (i.e., the asymmetry) appear to be reflected in the empirical literature reviewed in Section 4.2, where the evidence for an angel stock valuation premium is mixed while the evidence for a sin stock discounts appears consistent.

D. Subjective Probabilities

We now turn to assessing the effect negative and positive externalities have on subjective probability estimates. The histogram in Figure 3 compares true (objective) versus subjective (reported) probabilities for various true probability bins. The Bayesian objective probability of the stock being of the high return type is calculated as follows: given a history of n doubling and m halving outcomes, it is $2^{n-m}/(2^{n-m}+1)$. Given the discrete nature of signals and the asymmetry of updating, in most rounds, the true probability will be one of 1/3, 1/2, and 2/3 (corresponding to $|n-m| \leq 1$). We therefore use five true probability bins: below 1/3, exactly 1/3, exactly 1/2, exactly 2/3, and above 2/3. The thick black bars indicate the expected true probability for the corresponding bin (exactly 1/3, 1/2, or 2/3 for the inner bins, and the expected value of the true probabilities in the outer bins). The bars depict the subjective probability for the different treatments.

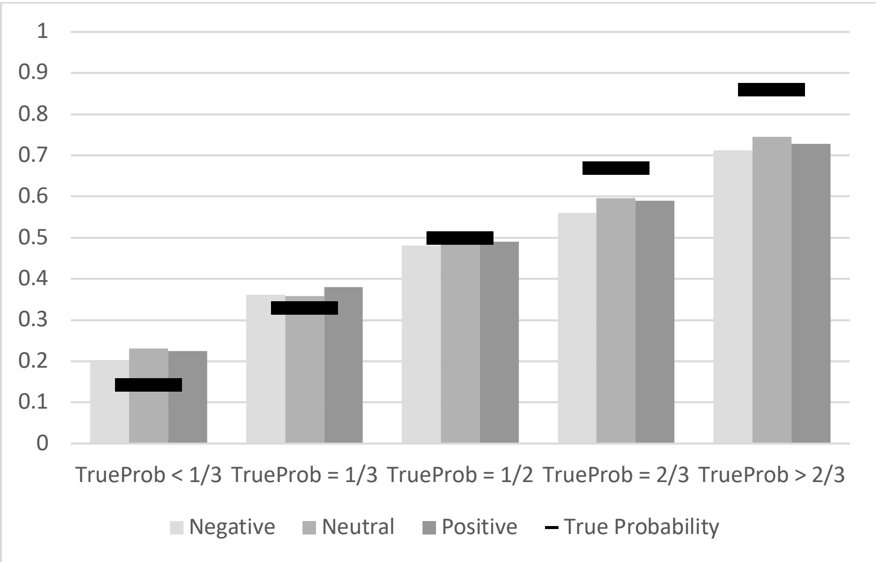

Figure 3: Objective and Subjective Probabilities
The figure depicts objective probabilities of the stock's type to subjective (reported) probabilities. The thick black bars indicate the expected true probability for the corresponding bin (exactly 1/3, 1/2, or 2/3 for the inner bins, and the expected value of the true probabilities in the outer bins). The bars depict the average subjective (reported) probability for the different treatments.

Two important patterns are clear from Figure 3. First, consistent with a large prior literature (e.g., Tversky and Kahneman, 1992, Abdellaoui, et. al., 2011, and Kuhnen, 2015), we find that subjective probabilities are "shrunk" toward the unconditional prior of 1/2. That is, when objective probabilities are less (more) than 1/2, subjects' perception of probabilities are too high (low). Second, we find that subjective probabilities in the Negative condition tend to be lower than the probabilities in the Positive condition. This difference is around 2 percentage points across objective probability bins.

The first pattern noted above resembles the behavior of "subjective weights" in non-expected utility models, such as Cumulative Prospect Theory. In particular, Prelec (1998) suggests the following formulation, adapted to our multi-treatment setting:

$$Subjective\ weight = \exp\left(-\left(\delta + \delta_{Pos} + \delta_{Neg}\right) * \left(-\log(True\ Prob)\right)^{\gamma+\gamma_{Pos}+\gamma_{Neg}}\right). \quad (1)$$

In the specification above, the Neutral treatment is the baseline. One can roughly think of the δ's as level parameters, shifting subjective probabilities up or down relative to objective probabilities, and of the γ's as curvature parameters. The Bayesian null is $\delta = \gamma = 1$ and $\delta_{Pos} = \delta_{Neg} = \gamma_{Pos} = \gamma_{Neg} = 0$. We adopt this formulation and fit it to the participants' stated beliefs, across all rounds.

Variable	Estimate	s.e.
δ	0.9282	0.0113
γ	0.6329	0.0170
δ_{Neg}	0.04618	0.0164
γ_{Neg}	-0.0286	0.0238
δ_{Pos}	-0.0022	0.0159
γ_{Pos}	-0.0267	0.0237

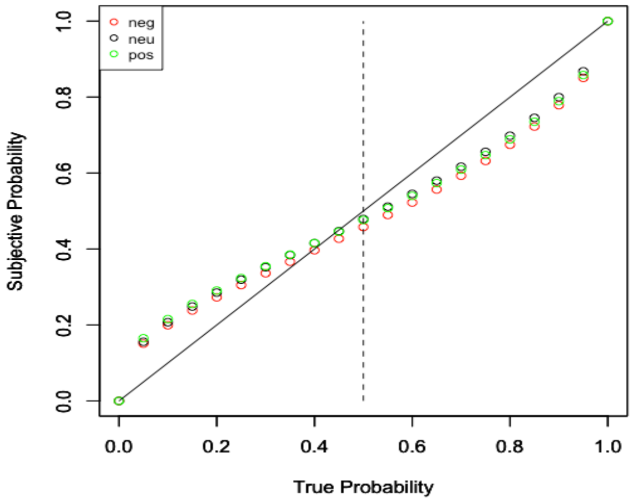

Figure 4: Cumulative Probability Function
This figure shows the maximum-likelihood estimation of equation (1) along with a plot that depicts the estimated parameters. Standard errors are in parentheses.

We report in Figure 4 the estimated subjective probability function parameters across the three treatments and plot the fitted treatment-dependent subjective probability as a function of objective probabilities. The subjective probability plots in the Positive and Neutral treatments are very close, while those of the Negative treatment are consistently below the other two, across the entire range of objective probabilities. The shift down, in the Negative treatment, appears to be similar across the range, consistent with the treatment effect on curvature being insignificant.

The asymmetry between the Negative and Positive treatments is borne out by the estimation results. There appears to be no statistically significant difference between the Positive and Neutral ones. However, both δ_{Neg} and γ_{Neg} are different from the corresponding Neutral and Positive treatment coefficients.

In Appendix D we employ the Mobius (2022) model to pin down the source of asymmetry between updating in the Negative versus the Positive/Neutral treatments. The analysis suggests that the probability assessment bias observed in the Negative treatment arises from lower sensitivity to "good news" about the stock's payoff distribution.

All in all, the belief results show modest, but surprising effects. Overall, we find no effect on beliefs or their updating in the Positive treatment. In contrast, we find muted response to positive information in the Negative treatment that is consistent with lower subjective probability estimates in this treatment over the high objective probability range.

Finding an effect on beliefs is surprising given that reported subjective probabilities, unlike stock allocations, have no externality on the linked cause; in fact, subjects are explicitly paid for accuracy. Second, the experiment is set up in a way that delinks the prospects of the stock from the cause and so beliefs should have remained constant even if the linked cause is one over which the subject has strong preferences. That said, our results on beliefs and updating echo the findings around at least two well-established behavioral biases – the halo effect and the confirmation bias. The halo effect suggests that positive or negative impressions carry over irrationally across domains in a way that is often unconscious and that can be affected by ESG policies (e.g., Nissbet and Wilson (1977), Hong et al (2019)). The confirmation bias suggests that individuals overweight evidence that is consistent with their priors (see Klayman (2008) for a survey). While, again, our findings are consistent with these biases, our contribution is to show that they are invoked only by the concerns for negative, but not for positive, externalities. The spillover effect of framing on beliefs is also consistent with Kuhnen (2015) who finds that subjects' learning process – in this very task without the externality treatments – differs based on whether outcomes are in the domain of gains or losses.

IV. Further Discussion

Our main findings concern the asymmetric impact of negative RI externalities on investment preferences and belief formation, and on the distribution of sensitivities to RI externalities in our sample of subjects. These findings, in turn, separately contribute to different branches of the RI literature.

A. Pecuniary Alignment

Because ESG rankings may correlate with perceived risks, such as future regulations (Krueger, Sautner and Starks, 2020; Stroebel and Wurgler, 2022), the question of whether RI represents an important departure from the traditional paradigm of balancing risk-return financial tradeoffs has been hotly debated. Investors who associate high ESG-rated stocks with higher expected returns or lower risk will allocate more capital to these investments, much as would investors who are motivated primarily by non-pecuniary motives.[18] While a number of studies attempt to establish non-pecuniary motives as an important aggregate investment channel, this is challenging to confirm in practice. Because we have designed our experiment to restrict variation to non-pecuniary externalities across the three treatments, we can directly observe the presence of non-pecuniary preferences. The results show that non-pecuniary

[18] For discussions of the different types of RI investment motivations, see, for example, Edmans and Kacperczyk (2022) and Starks (2023).

preferences are pervasive across subjects and economically large, which supports the broader claims of empirical studies pointing to the importance of such preferences and where controlling for financial outcomes is difficult or impossible (see references in Footnote 4).

B. Asymmetry in RI Approaches and Impact on Stock Returns

Our findings resonate with a number of empirical stylized facts, suggesting they may help provide the basis for microfounding a descriptive model of individual RI preferences that can also explain aggregate behavior. For example, our asymmetric results on preferences in the presence of negative versus positive externalities are consistent with the asymmetry found in studies of fund flows and stock market reactions which appear to be more strongly affected by negative RI events.[19] Similar to our results but under a different experimental setting, Chew and Li (2021) find strong evidence for asymmetry in sin stock aversion versus virtue stock affinity, which they model by appealing to the literature on source-dependent risk aversion.

Our results also conform to what we observe in practice in the RI market. Although much of RI has evolved to include positive ESG tilting, negative screening remains a pervasive strategy.[20] An often-followed model of this investment approach is the Norwegian sovereign wealth fund (the Government Pension Fund Global) which invests sustainably but also excludes companies that they believe do not meet their ethical norms.[21]

Our experimental design allows us to control for the magnitude of the externality while switching its sign. This is not generally possible in natural experiments, for example, where ordinal ESG rankings are used to compare RI alternatives. In addition to providing a "clean"

[19] See, for example, Krueger (2015); Bialkowski and Starks (2018); Hartzmark and Sussman (2019).
[20] For example, of the 86 U.S. investment managers who reported their screening technologies to the Principles of Responsible Investing, 91% use some type of negative screening strategy. In fact, 33% use only negative screening while 58% use negative screening combined with some form of positive screening. Similarly, a survey finds that 69% of managers use negative screening in their investment decisions (US SIF, 2020). We thank the PRI for providing us with 2018 data to calculate these statistics. The PRI has not reviewed the methodology applied to, use of, or conclusions drawn from this data.
[21] https://www.nbim.no/en/the-fund/responsible-investment/.

causal channel for observed asymmetry, our experiment allows us to rule out the possibility that the asymmetry exists because poorly ranked ESG alternatives are significantly worse from a measurable social perspective than highly ranked ESG alternatives.

Furthermore, RI theoretical models in which investors exhibit a pronounced preference for stocks that rate highly on an E, S or G dimension predict a valuation premium for such stocks (e.g., Pastor, Stambaugh, and Taylor, 2021). The empirical literature, however, lacks a consensus around this intuitive prediction.[22] Our experimental findings of only a weak preference for investing in stocks linked to positive externalities is consistent with the weaker evidence in the empirical literature for the performance of strategies that favor investment in ESG. Our stronger findings concerning negative externalities are consistent with the clearer empirical evidence of shunned stock undervaluation (e.g., Hong and Kacperczyk, 2009; Statman and Glushkov, 2009). Thus, our findings provide a way to interpret the array of evidence documented in a large and growing literature and should help in the development of new theories.

C. The Role of Distorted Beliefs

Importantly, our study points to another potential (albeit weaker) channel impacting asset markets by RI investors. Belief formation impacts trading (see, for instance, Füllbrunn et al., 2022). If the ability of individuals to infer the likelihood of outcomes is impacted by negative externalities as our results imply, asset prices will be affected beyond what is suggested in earlier theoretical work that focuses only on the impact of tastes on allocations. If a sufficiently large proportion of RI investors in the economy deviates from Bayesian updating because of RI sensitivities, then additional distortions could arise from the effective presence

[22] See, for example, Statman and Glushkov (2009), Edmans (2011), Humphrey, Lee, and Shen (2012), Bansal, Wu, and Yaron (2022), Hwang, Titman, and Wang (2022) and Liang, Sun, and Teo (2022)

of "pessimistic" investors. Because not all models incorporate non-Bayesian beliefs, it seems useful to disentangle the impact of RI on preferences versus beliefs. Our results provide empirical guidance on how this can be done.

D. Distribution of RI Preferences and Theoretical RI Models

Our results, particularly the results regarding the effects of negative RI information on investors' asset allocations and beliefs, have important implications for financial markets. If the percentage of RI investors grows in the economy as still expected by many, asset prices are likely to be affected by their allocation choices, which has been shown theoretically and suggested empirically. For example, Heinkel, Kraus and Zechner (2001) and Luo and Balvers (2017) demonstrate how shunning sin stocks would have the effect of driving these stocks' prices lower, which has empirical support in Hong and Kacperczyk (2009), Statman and Glushkov (2009), as well as Chava (2014). Similarly, Fama and French (2007) provide a simple theoretical framework to demonstrate that investor tastes, such as tastes for responsible investing, can distort pricing in asset markets. They show that these distortions in prices could be large under certain circumstances: when investors with particular tastes represent a substantial fraction of invested wealth; when the investors have such tastes for a wide range of assets; when investors' positions vary quite a bit from the market portfolio; and when the returns on the investors' underweighted assets are not highly correlated with the returns on their over-weighted assets. In other words, it is plausible to expect an impact on asset prices when responsible investors represent a substantial percentage of investors in the market.

As many as half of our subject pool would substantially reduce their allocation to an asset with a negative RI association. The magnitude of such RI preferences, when extrapolated to the population, should be sufficient to significantly impact asset prices. Moreover, the asymmetry we document in RI preferences is consistent with only a subset of existing theoretical models (as we point out in Footnote 4).

E. Additional Experimental Studies

Three recent studies, all of which are based on Dutch investors, examine investors' choices of responsible investing strategies. Riedl and Smeets (2017) conclude that intrinsic social preferences and social signaling are the investors' primary motivations and that while financial motivations enter into the decision making, they play a relatively minor role. Brodback, Guenster and Mezger (2019) find a positive link between altruistic values and the relative importance of social responsibility to investors. They also find that the link strengthens under certain conditions: when individuals believe their investments can make a social or environmental impact or when they feel moral obligations regarding their investments. Lastly, Bauer, Rouf and Smeets (2021) conduct a field experiment in which Dutch pension participants are allowed to vote on whether the pension system should follow three or four of the United Nations Sustainable Development Goals. They conclude that the choice of 66% of the participants to follow more of the goals, i.e., engage in more responsible investing activities, is based on nonfinancial rather than financial considerations. Whereas these papers seek to answer the question of *why* investors select into being RI investors, our work investigates *how* the RI information is incorporated into investors' decisions.

Our analysis is also complementary to several contemporaneous experimental studies, but the focus and consequently, the experimental design, exhibit key differences. Bonnefon et al. (2022) examine the private valuation assigned by MTurk subjects to direct giving to (or taking from) charities, and find it is roughly linear in the small stakes considered. Moreover, private valuations do not significantly depend on whether a subject is pivotal to the giving (i.e., whether the amount the charity receives depends on actions taken by the subject). Our design is fundamentally different in that it incorporates and examines dimensions of quantity, uncertainty, and learning linked to the RI decision with a specific focus on asymmetry. On the other hand, we do not test for a difference between pivotal and non-pivotal treatments. Whereas both our paper and Bonnefon et al. provide strong evidence that tastes matter in evaluating RI, their paper focuses on the difference in pivotal vs. non-pivotal effects and appears to find no evidence for the strong asymmetric results discovered in our setting.[23] In their paper, there is no room for updating so they cannot separately assess the effect of the treatment on beliefs vs. allocations.

Brodback, Guenster and Pouget (2020) also conduct an experiment with charitable donations tied to a financial investment, in this case an initial public offering. They find that individuals have a price premium for social responsibility that also depends on the asset's financial performance. Their focus is on the pricing of the social benefit in combination with the financial performance while our focus is on how the externalities affect beliefs and allocations. Likewise, Heeb et. al. (2021) examine investors' willingness to pay for sustainable investment in an experimental setting. They focus on positive externalities and how willingness-to-pay changes with impact and the choice set.

Finally, we point out that a vast literature exists on "other-regarding behavior," mostly focused on strategic choice problems.[24] Although gain-loss asymmetry, introduced in Kahneman and Tversky (1979), is one of the most influential and persistent stylized facts in human decision making, the evidence we find for its social preference manifestation appears to be new.

[23] The significantly smaller stakes employed in Bonnefon et al. (2022) for both subjects and charities might serve to mask a difference between the pivotal and non-pivotal treatments, or an asymmetry between the impact of negative versus positive payoffs to the charity.
[24] See Cooper and Kagel (2016) for a review.

V. Conclusions

In this paper we employ an experimental setting to study how social externalities influence individuals' investment decisions and belief formations. We find that linking individual investment outcomes to social externalities significantly affects individuals' allocations between a risky asset and cash as well as their subjective beliefs over the investment outcomes.

Most importantly, due to our unique experimental design we are able to demonstrate that asymmetry is large, pervasive, and fundamental to non-pecuniary investor preferences. In particular, we estimate the strength and prevalence of RI influences on investment across subjects and find that roughly a third to half of our subjects would reduce their allocation to a lucrative risky investment by an average of nearly 50% if that investment were linked to negative social externalities. In contrast, the impact of positive externalities on investment decisions pales by comparison. We also find asymmetry in belief formations in that subjects are significantly more pessimistic about investment outcomes when the investment is linked to a negative externality, even when controlling for the objective prospects of investments, while positive externalities do not appear to influence beliefs. The type, magnitude, and pervasiveness of the effects we identify can be readily incorporated into the modeling, and therefore equilibrium consequences, of RI preferences.

In magnitude, the asymmetries we find rival the strength of loss aversion estimates in the literature. Interestingly, evidence of charitable-giving motives appears to be largely absent in the data. The effect we find is not only high in magnitude, but also pervasive among a large portion of the subject pool. This suggests that, based on existing theoretical work, when extrapolated to a market setting our findings would translate into observable price impacts, but only for firms associated with negative externalities. Indeed, a survey of the empirical RI literature suggests not only a prevalence of negative screening – an important component of the majority of RI strategies – but that only "sin stocks" consistently exhibit greater discounting than warranted based on conventional risk adjustment (evidence for the overpricing of "angel stocks" tends to be mixed). Another of our novel experimental findings is that social preferences affect investors' subjective probabilities about their investments. Although this latter effect is modest, it reflects the importance that social preferences can have on how investors process information (e.g., update their beliefs).

Responsible investing has become an increasingly important aspect of some individuals' investment opportunity sets. Theory and empirical evidence have demonstrated that growing tastes for responsible investing can impact asset pricing. Our findings help refine existing facts and insights by pointing to novel drivers of responsible investment. Importantly, our results have implications for how to think about incorporating social preferences into theoretical models because they demonstrate not only how preferences and beliefs are affected but also the large heterogeneity that exists in these effects. In addition, our results have implications for policy. In particular, the strong asymmetric effects we find suggest that, from an investor's perspective, the marginal benefit of reducing harm is much greater than the marginal benefit of doing good.

References

Abdellaoui, M., Baillon, A., Placido, L., and Wakker, P. "The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation." *American Economic Review,* 101 (2011), 695-723.

Abdellaoui, M., Bleichrodt, H., and Paraschiv, C. "Loss Aversion under Prospect Theory: A Parameter-Free Measurement." *Management Science,* 53 (2007), 1659-1674.

Bansal, R., Wu, D., and Yaron, A. "Socially Responsible Investing in Good and Bad Times," *Review of Financial Studies,* 35 (2022), 2067-2099.

Barber, B., Morse, A., and Yasuda, A. "Impact Investing," *Journal of Financial Economics*, 139 (2021), 162-185.

Bauer, R., Ruof, T., and Smeets, P. "Get Real! Individuals Prefer More Sustainable Investments." *Review of Financial Studies,* 34 (2021), 3976-4043.

Benjamin, D. J. "Errors in Probabilistic Reasoning and Judgment Biases." *Handbook of Behavioral Economics: Applications and Foundations 1*, 2 (2019), 69-186.

Bialkowski, J., and Starks, L.T. "SRI Funds: Investor Demand, Exogenous Shocks and ESG Profiles." Working paper (2018).

Bolton, P. and Kacperczyk, M. "Do Investors Care About Carbon Risk?" *Journal of Financial Economics,* 142 (2021), 517-549.

Bonnefon, J-F., Landier, A. , Sastry, P., and Thesmar, D. "The Moral Preferences of Investors: Experimental Evidence." NBER Working Paper 29647 (2022).

Brodback, D., Guenster, N., and Mezger, D. "Altruism and Egoism in Investment Decisions. *Review of Financial Economics,* 37 (2019), 118-148.

Brodback, D., Guenster, N., and Pouget, S. "The Valuation of Corporate Social Responsibility: A Willingness-To-Pay Experiment." Working paper (2020).

Chava, S. "Environmental Externalities and Cost of Capital." *Management Science,* 60 (2014), 2223-2247.

Chew, S.H., and Li, K.K. "The Moral Investor: Sin Stock Aversion and Virtue Stock Affinity." Working Paper (2021). *Available at SSRN 3773971*.

Chowdhry, B., Davies, S.W., and Waters, B. "Investing for Impact." *Review of Financial Studies,* 32 (2019), 864-904.

Cooper, D. J., and Kagel, J. "Other-regarding Preferences." *The Handbook of Experimental Economics* (2016), 217-289.

Dimson, E., Marsh, P., and Staunton, M. "Exclusionary Screening." *The Journal of Impact and ESG Investing* (2020), 66-75.

Edmans, A. “Does the Stock Market Fully Value Intangibles? Employee Satisfaction and Equity Prices.” *Journal of Financial Economics,* 101 (2011), 621-640.

Edmans, A., and Kacperczyk, M. “Sustainable Finance.” *Review of Finance,* 26 (2022), 1309-1313.

Engel, C. “Dictator Games: A Meta Study.” *Experimental Economics* 14 (2011), 583-610.

Fama, E., and French, K. “Disagreement, Tastes and Asset Prices” *Journal of Financial Economics* 83 (2007), 667-689.

Flammer, C. “Corporate Green Bonds.” *Journal of Financial Economics* 142 (2021), 499-516.

Füllbrunn, S., Huber, C., Eckel, C., and Weitzel, U. “Heterogeneity of Beliefs and Trading Behavior: A Reexamination.” *Journal of Financial and Quantitative Analysis,* 57 (2022), 1-25.

Geczy, C. C., Jeffers, J.S., Musto, D.K. and Tucker, A.M.. “Contracts With (Social) Benefits: The Implementation of Impact Investing.” *Journal of Financial Economics,* 142 (2021), 697-718.

Goldstein, I., Kopytov, A., Shen, L., and Xiang, H.. “On ESG Investing: Heterogeneous Preferences, Information and Asset Prices.” Working Paper (2023).

Hartzmark, S. M. and Sussman, A.B.. “Do Investors Value Sustainability? A Natural Experiment Examining Ranking and Fund Flows.” *Journal of Finance,* 74 (2019), 2789-2837.

Heeb, F., Kölbel, J., Paetzold, F., and Zeisberger, S.. “Do Investors Care About Impact?” Working paper (2021).

Heinkel, R., Kraus, A. ,and Zechner, J. “The Effect of Green Investment on Corporate Behavior.” *Journal of Financial and Quantitative Analysis,* 36 (2001), 431-449.

Hong, H. and Kacperczyk, M. “The Price of Sin: The Effects of Social Norms on Markets.” *Journal of Financial Economics,* 93 (2009), 15-36.

Humphrey, J. E., Lee, D.D., and Shen, Y. “Does it Cost to be Sustainable?” *Journal of Corporate Finance,* 18 (2012), 626-639.

Hwang, C.Y., Titman, S., and Wang, Y. “Investor Tastes, Corporate Behavior, and Stock Returns: An Analysis of Corporate Social Responsibility” *Management Science,* 68 (2022), 7065-7791.

Kahneman, D., Knetsch, J.L., and Thaler, R.H. “Experimental Tests of The Endowment Effect and The Coase Theorem.” *Journal of Political Economy,* 98 (1990), 1325-1348.

Kahneman, D., and Tversky, A. “Prospect Theory: An Analysis of Decision Under Risk.” *Econometrica,* 47 (1979), 263-291.

Kensinger, E. A., and Schacter, D.L. “Memory and Emotion.” *Handbook of Emotions,* 3 (2008), 601-617.

Kuhnen, C. “Asymmetric Learning from Financial Information.” *Journal of Financ,e* 70 (2015), 2029-2062.

Krueger, P. “Corporate Goodness and Shareholder Wealth.” *Journal of Financial Economics,* 115 (2015), 304-329.

Krueger, P., Sautner, Z. and Starks, L.T. “The Importance of Climate Risk For Institutional Investors.” *Review of Financial Studies,* 33 (2020), 1067-1111.

Larcker, D.F. and Watts, E.M. “Where’s the Greenium?” *Journal of Accounting and Economics,* 69 (2020), Article 101312.

Liang, H., Sun, L., and Teo, M. “Responsible Hedge Funds” *Review of Finance,* 26 (2022), 1585-1633.

List, J. “On the Interpretation of Giving in Dictator Games” *Journal of Political Economy,* 115 (2007), 482-493.

Luo, H.A. and Balvers, R. “Social Screens and Systematic Investor Boycott Risk.” *Journal of Financial and Quantitative Analysis,* 52 (2017), 365-399.

Mellers, B., Schwartz, A. and Ritov, I.. “Emotion-based Choice.” *Journal of Experimental Psychology: General,* 128 (1999), 332-345.

Möbius, Markus M., Niederle, M., Niehaus, P. and Rosenblat, T.S. “Managing Self-Confidence: Theory and Experimental Evidence.” *Management Science,* 68, (2022), 7793-7817

Oehmke, M., and Opp, M. “A Theory of Socially Responsible Investment.” Working Paper (2023).

Pastor, L., Stambaugh, R. and Taylor, L. “Sustainable Investing in Equilibrium.” *Journal of Financial Economics,* 142 (2021), 550-571.

Pedersen, L.H., Fitzgibbons, S. and Pomorski, L. “Responsible Investing: The ESG-Efficient Frontier.” *Journal of Financial Economics* 142 (2021), 572-597.

Prelec, D. “The Probability Weighting Function.” *Econometrica,* 66 (1998), 497-528.

Riedl, A. and Smeets, P. “Why Do Investors Hold Socially Responsible Mutual Funds?” *Journal of Finance,* 72 (2017), 2505-2550.

Slovic, P., and Lichtenstein, S. “Comparison of Bayesian and Regression Approaches to the Study of Information Processing in Judgment, *Organizational Behavior and Human Performance,* 6 (1971), 649-744.

Starks, L.T. “Presidential Address: Sustainable Finance and ESG Issues – *Value* Versus *Values*.” *Journal of Finance,* 78 (2023), 1837-1872.

Statman, M., and Glushkov, D. “The Wages of Social Responsibility.” *Financial Analysts Journal,* 65 (2009), 33-46.

Stroebel, J., and Wurgler, J. “What Do You Think About Climate Finance?” *Journal of Financial Economics,* 142 (2021), 487-498.

Svenson, O. “Are We All Less Risky and More Skillful Than Our Fellow Drivers?” *Acta Psychologica*, 47 (1981), 143-148.

Tversky, A. “On the Psychology of Prediction.” *Psychological Review,* 80 (1973), 237-251.

Tversky, A., and Kahneman, D. “Advances in Prospect Theory: Cumulative Representation of Uncertainty.” *Journal of Risk and Uncertainty,* 5 (1992), 297-323.

US SIF. Report on US Sustainable, Responsible and Impact Investing Trends 2020, Washington, DC: US SIF. (2020).

Yoo, P. “ESG Investing: A Tale of Two Preferences.” Working Paper (2022).

Zerbib, O.D. “A Sustainable Capital Asset Pricing Model (S-CAPM): Evidence from Environmental Integration and Sin Stock Exclusion.” *Review of Finance,* 26 (2022), 1345-1388.

Appendix A: Experiment Snapshots

Panel 1: Initial Round Allocation Screen

First Block, Trial One

You are about to begin a new trial consisting of six investing rounds.

The **computer will now randomly select the type of** stock you will be facing during this trial.
There is a **50% chance that the computer will select a high payoff stock** and a 50% that it will select a low payoff stock.

A **high payoff doubles with probability 2/3** and halves with probability 1/3.
A **low payoff stock doubles with probability 1/3** and halves with probability 2/3.

Estimate the probability that this is the **high payoff** stock.

Please enter a number between 0 and 100:

You have **100 ECU**. Enter the amount you wish to invest in the stock in the next round. The rest will be invested in risk-free cash.

Stock:

Panel 2: Outcome Screen after Three Rounds

Trial: 1

The stock outcome for Round 3 is: Halve (↓)

History of Decisions and Outcomes

Round Number	Probability Estimate Prior to Outcome	Stock Allocation Prior to Outcome	Outcome	Winnings from Round
1	20	50	Halve (↓)	75
2	30	50	Double (↑)	150
3	40	40	Halve (↓)	80
4				
5				
6				

Recall that **high payoff** stocks will double with a probability of 2/3 and halve with a probability of 1/3.

Low payoff stocks will double with a probability of 1/3 and halve with a probability of 2/3.

Estimate the probability that this is the **high payoff** stock.

Please enter a number between 0 and 100:

You have **100 ECU**. Enter the amount you wish to invest in the stock in the next round. The rest will be invested in risk-free cash.

Stock:

Panel 3: Social Issues

What Social Issues Are Important to You?

Below are a number of current corporate social issues, listed in alphabetical order. Please rank these issues from 1 to 6 in terms of how much you care about each issue. A ranking of 1 means you care the most and a ranking of 6 means you care the least. You can rank by clicking and dragging. If you would like to keep the current ranking simply click and slightly drag any of the options.

Gender discrimination (example: not hiring women for jobs for which they are qualified)

Human Trafficking (example: forcing children to work as slave labour in factories or on farms)

Refugees (example: trading with dictatorial regimes which have resulted in mass fleeing from the country)

Poverty (example: underpaying workers in developing countries)

Animal welfare (example: infecting laboratory monkeys with diseases to test pharmaceutical drugs)

Environment (example: releasing carbon dioxide into the atmosphere)

What Social Issues Are Important to You?

You selected environment as your top social issue. Below is a description of two environmental non-profits.

The Rainforest Alliance works to conserve biodiversity and ensure sustainable livelihoods by transforming land-use practices, business practices and consumer behaviour.

Conservation Strategy Fund sustains natural ecosystems and human communities through strategies powered by conservation economics. Our trainings, analyses and timely expertise make development smarter, quantify the benefits of nature, and create enduring incentives for conservation.

Please select a non-profit organization to link to your trading profits:

- ○ The Rainforest Alliance
- ○ The Conservation Strategy Fund

Panel IV: Positive Block Instruction Screen and Outcome Screen

The Conservation Strategy Fund gains by your stock investment

You are about to begin a block of two trials. Each trial consists of six investment rounds. You will be asked to make the same decision as in the practice trial.

At the beginning of each of the two trials the computer will randomly choose whether the stock for that trial is a high or low payoff stock. You will then face the series of six investment rounds.

In each investment round you will be asked how much you wish to allocate to the stock and to cash. After the stock outcome is determined, you will be asked to indicate the probability that you are facing a high payoff stock and how much you trust your probability estimate.

One round may be randomly chosen as your actual payoff round. If that round is chosen as your payoff round, in addition to paying you, **the researchers will DONATE the same amount of money as your payoff from your stock investment to your designated non-profit**. For example, if you earn 135 ECU from investing including 70 ECU from your stock investment, your chosen non-profit will receive 70 ECU.

The payment to the non-profit will come out of the researchers' funding, not from your trading payoff. You will be sent an email documenting our payments to the various non-profit organizations involved in this experiment in a few days.

Your designated non-profit for this block of three trials is: The Conservation Strategy Fund

Non-profit gains by your stock investment

Trial: 1

The stock outcome for Round 3 is: Double (↑)

History of Decisions and Outcomes

Round Number	Probability Estimate Prior to Outcome	Stock Allocation Prior to Outcome	Outcome	Winnings from Round (if round is selected)	Payment to The Conservation Strategy Fund (if round is selected)
1	40	60	Double (↑)	160	120
2	60	45	Halve (↓)	77.5	22.5
3	50	50	Double (↑)	150	100
4					
5					
6					

Recall that **high payoff** stocks will double with a probability of 2/3 and halve with a probability of 1/3.

Low payoff stocks will double with a probability of 1/3 and halve with a probability of 2/3.

Estimate the probability that this is the **high payoff** stock.

Please enter a number between 0 and 100:

You have **100 ECU**. Enter the amount you wish to invest in the stock in the next round. The rest will be invested in risk-free cash.

Stock:

Panel 5: Negative Block Instruction Screen and Outcome Screen

Block: SPCA International loses by your stock investment

You are about to begin a block of two trials. Each trial consists of six investment rounds. You will be asked to make the same decision as in the practice trial.

At the beginning of each of the two trials the computer will randomly choose whether the stock for that trial is a high or low payoff stock. You will then face the series of six investment rounds.

In each investment round you will be asked how much you wish to allocate to the stock and to cash. After the stock outcome is determined, you will be asked to indicate the probability that you are facing a high payoff stock and how much you trust your probability estimate.

One round may be randomly chosen as your actual payoff round. If that round is chosen as your payoff round, in addition to paying you, **the researchers will DEDUCT from funds allocated to the non-profit the same amount of money as your payoff from your stock investment**. For example, if you earn 135 ECU from investing including 70 ECU from your stock investment, your chosen non-profit will have 70 ECU deducted from their balance.

The deduction from the non-profit will have no effect on your trading payoff. You will be sent an email documenting our payments to the various non-profit organizations involved in this experiment in a few days.

Your designated non-profit for this block of three trials is: SPCA International

Block: Non-profit loses by your stock investment

Trial: 1

The stock outcome for Round 3 is: Double (↑)

History of Decisions and Outcomes

Round Number	Probability Estimate Prior to Outcome	Stock Allocation Prior to Outcome	Outcome	Winnings from Round (if this round is selected	Deduction from SPCA International (if this round is selected)
1	50	30	Halve (↓)	85	-15
2	20	40	Halve (↓)	80	-20
3	15	30	Double (↑)	130	-60
4					
5					
6					

Recall that **high payoff** stocks will double with a probability of 2/3 and halve with a probability of 1/3.

Low payoff stocks will double with a probability of 1/3 and halve with a probability of 2/3.

Estimate the probability that this is the **high payoff** stock.

Please enter a number between 0 and 100:

You have **100 ECU**. Enter the amount you wish to invest in the stock in the next round. The rest will be invested in risk-free cash.

Stock:

Appendix B: Survey to Establish External Validity

Below is an example of a survey question eliciting subjects' preference for "reducing good" versus "causing harm". The full survey is available at https://uqbel.az1.qualtrics.com/jfe/form/SV_3VID8WMyu4q05z8.

A local town council debates the use of municipal funds to house and support people experiencing homelessness. Because of the limited funds available for all of the city needs, the council is evenly split on the measure, and the final vote resides with the elected mayor who is undecided (but will have to decide one way or another in a week). It is impossible to know which way the Mayor currently leans.

If you had to make a choice between the following alternatives, which would you choose

- **Donate** $1,000 to a firm that will lobby the Mayor to vote **against** helping people experiencing homelessness.
- **Cut** $1,000 of funding from a charity that will lobby the Mayor to vote **for** helping people experiencing homelessness.
- The choices are equally bad.

Appendix C: Closer Examination of Subject Heterogeneity

In a naïve approach we would regress allocations on subjective probabilities for each subject and treatment, but this presents two major challenges. First, a subject-by-subject regression analysis lacks power because we have only twelve decisions for each subject per treatment. Second, subject responses reflect experimental error for a variety of reasons.[25] Because allocations are bounded between 0 and 100 ECU, experimental noise can bias observations away from the bounds and correspondingly bias inferences about individual allocation responses to perceived probabilities.[26] To address these issues, we adopt a reduced form model relating noisy optimal allocations to subjective probabilities. We then estimate individual-level parameters as random effects. This allows us to efficiently quantify individual preferences as well as attempt to control for edge biases arising from noise.

We begin by noting that most utility models predict an optimal allocation that roughly resembles a sigmoid function of subjective probability.[27] Consistent with that, we assume that noisy individual allocations, in a given treatment, can be described in reduced form as follows:

$$ECU_i^T = a_i + b_i f(Prob_i) + e_i, \qquad (C1)$$

where the transformed allocation, $ECU^T \in \Re$, is unbounded and related to the allocation choice via the sigmoid transformation,

$$ECU = 100/(1 + \exp(-ECU^T)). \qquad (C2)$$

In Equation (C1), i refers to the subject, $f()$ is an increasing function to be specified soon, a_i and b_i are subject-specific constants, and e_i is experimental noise. A standard sigmoid

[25] Subjects make mistakes, get confused, or simply find it difficult to maintain the same level of engagement throughout the rounds of an experiment. There is also the possibility that they enjoy injecting a random component to their choices. Holding constant the subjective probability and treatment, the average filtered subject exhibited a standard deviation of 11 ECU in their allocations.

[26] While this bias can also impact the estimates in Table 2, within-subject estimates are more sensitive to the noise-induced bias. Indeed, consistent with a bias-driven deviation from the regression model, the residuals in Regression (2) of Table 2 exhibit a U-shape magnitude with respect to the perceived probability (roughly highest at the boundaries).

[27] A sigmoid function takes the form, $y = A\,(1 + \exp(a - bx))^{-1}$ with $b \geq 0$.

corresponds to the case where f is linear. Because ECU is convex in ECU^T for low values and concave for high values, this formulation models how experimental noise biases allocations away from the boundaries of [0,100].[28] The choice of $f()$ reflects behavior in the absence of noise. For instance, if $f(p)$ is finite at $p=0$, then the decision maker would be risk-loving because the allocation to the stock (an actuarially fair gamble at p=0) would be greater than zero. Because b is a measure of the sensitivity of allocation to a first-degree stochastically dominating shift in the payoff distribution of the stock, one can interpret b>0 as a measure of local risk tolerance with higher b signifying higher risk tolerance. We model

$$f(p) = \ln\left(\frac{p}{1-p}\right). \qquad (C3)$$

Our modeling choice implies global risk aversion and that, absent noise, as subjective probability approaches one the allocation will tend to the full endowment of 100 ECU. This lends parsimony to modeling the dependence of allocation on subjective probability.[29]

To address the issue of statistical power, we adopt a random effects framework estimated to allow for treatment differences. That is, we allow each subject to have a different average level of investment (intercept, a_i) and a different allocation sensitivity to probabilities (slope, b_i) in each of the treatments. In a random effects framework, the subject-level coefficients are assumed to be drawn from a distribution whose mean and standard deviation are estimated. A benefit of this approach over individual regressions is the joint estimation of covariance of distinct random effects. The estimated model is

$$ECU^T_{\tau,i,n} = a_{\tau,i} + b_{\tau,i} f\left(Prob_{\tau,i,n}\right) + e_{\tau,i,n}. \qquad (C4)$$

---

[28] This is an implication of Jensen's Inequality.

[29] To avoid infinities in expressions (2) and (3), extreme allocations of 0 or 100 ECU are adjusted to 0.1 or 99.9 ECU, respectively, and extreme reported probabilities of 0 or 1 are adjusted to 0.001 or 0.999. The qualitative conclusions of our analysis are robust to simply excluding these observations. They are also robust to alternative specifications such as, $f(p) = p$, or $f(p) = ln(p)$, or $f(p) = c\,ln(p) + dp$, or $f(p) = c\,ln(p) + ln(1-p)$.

where τ denotes the treatment, *i,* the subject ID (corresponding to the random effect), *n,* the round number for the given treatment, and $Prob_{\tau,i,n}$, the probability reported by the subject in round *n* of the given treatment.

Results are reported in Table C-1 and correspond to the estimated means, standard deviations, and correlations of the distribution of random coefficients.[30] Table C-1 also reports the residual standard deviation. Consistent with our prior results, we observe that in the Negative treatment, relative to the Neutral one, both the baseline allocation, mean(*a*), and the sensitivity to probabilities, mean(*b*), are significantly lower. By contrast, the two Positive treatment mean coefficients do not allow for a simple ranking relative to the Neutral treatment estimates (the joint difference is statistically insignificant with a *p* value of 36%).

Table C-1: Random Effects Regression of Transformed Allocation on Reported Probabilities

The table reports random-effect regressions of transformed allocations on reported probabilities estimated as specified in Equation (C4). *a* refers to intercept and *b* to sensitivity to probabilities. The error term is assumed to be i.i.d. across subjects and rounds, within treatment. Standard errors are in parentheses. Asterisks denote the probabilities for statistical tests of differences from zero (**p*<0.1; ***p*<0.05; ****p*<0.01).

Treatment	mean(*b*)	sd(*b*)	mean(*a*)	sd(*a*)
Negative	0.652***	0.699**	-1.733***	1.880***
	(0.082)	(0.081)	(0.177)	(0.132)
Neutral	0.966***	0.581***	-0.751***	1.651***
	(0.069)	(0.073)	(0.156)	(0.117)
Positive	0.903***	0.857***	-0.583***	1.690***
	(0.095)	(0.087)	(0.159)	(0.119)
	sd(e_{Neg})	sd(e_{Neu})	sd(*ePos*)	*N*
	1.904***	1.805***	1.684***	4500
	(0.038)	0.036	0.034	

Second, the table shows that the slope dispersion across subjects, sd(*b*), is estimated to be larger in the Negative and Positive treatments, compared with the Neutral treatment. This

[30] We only report significant random effect correlations.

suggests pronounced heterogeneity in treatment effects across subjects. Finally, we observe large and heterogeneous levels of residual variation across the treatments. The residual magnitudes justify our effort to adjust for edge-induced biases.[31]

Figure C-1, Panel A, plots the estimated allocation using the mean coefficients from Table C-1 and setting the experimental noise to zero. The plot includes 95% confidence interval bars for the Positive and Negative treatments.[32] From this figure we observe that the negative treatment effect is most pronounced when the subjective probability is higher than 50%. The figure is consistent with the absence of an average charitable-giving motive which, if present, would imply a significantly higher allocation in the Positive treatment at low probabilities. As with Figure 1, the Positive treatment allocations are generally above those of the Neutral treatment, but we stress that the two Neutral and Positive treatment curves cannot be differentiated statistically because they are generated by coefficients that are (jointly) not statistically different.

Figure C-1, Panel B, plots the magnitude of the deviations of each of the two treatments from the Neutral case. Holding constant the subjective probability, the magnitude of the externality is the same in the Negative and Positive treatment (with the opposite sign, of course). The plotted deviations from the Neutral treatment can therefore be viewed as responses to an equal-magnitude positive versus negative stimuli. From the plot, when the subjective probability is at 50%, the estimated model response is more than four times larger to a Negative "stimulus" than a positive one. This is consistent with the previous estimates and rivals loss-aversion in strength.

---

[31] The random effects model estimates correlations between the *a*'s and *b*'s, some of which are found to be statistically different from zero. In the interest of brevity, we omit these results.

[32] The Neutral treatment confidence intervals, like the Neutral treatment point estimates, are close to those of the Positive treatment. We omit these for visual clarity.

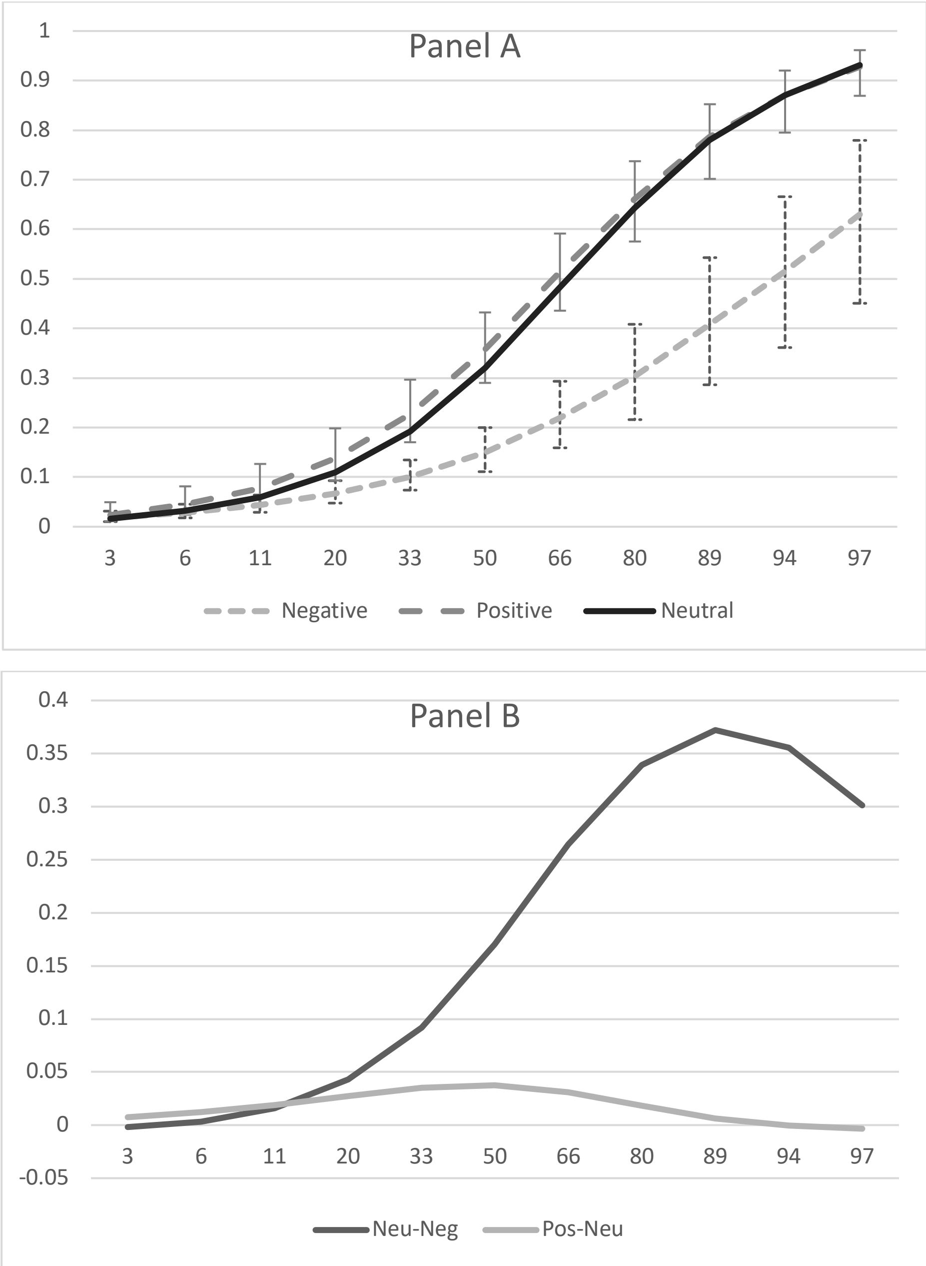

Figure C-1: Expected Allocations Across Treatments
This figure depicts estimated model predictions of the treatment effects on allocation (y-axis) against the subjective probability that the stock is of the high type (x-axis). The estimated model prediction sets experimental noise to zero. By contrast, Figure 1 is based on aggregated statistics that potentially include bias from experimental noise. Panel A reports the expected allocation across treatments while Panel B reports the expected allocations relative to the neutral treatment (i.e., neutral-negative and positive-neutral).

To compare model predictions to the simple statistics from Table 1, Panel A, we use the model estimates of each subject's intercept and slope parameters to calculate that subject's *expected allocations* in the treatment without experimental noise. We do this using the theoretical distribution of objective stock probabilities in the six rounds.[33]

Table C-2 reports the average over subjects of their predicted expected allocation in each of the treatments. Consistent with the results of Table 1, the expected allocations are lower in the Negative treatment (0.262) relative to the Neutral treatment (0.373) which, in turn, is not far from the expected allocation in the Positive treatment (0.403). The magnitude of the treatment effect is similar to that documented in Table 1 and corresponds to a reduction of nearly 30% in allocations in going from the Neutral to the Negative treatment.

Table C-2: Average of Expected Allocations in Each Treatment

The table reports the expected allocation averaged across subjects in each of the treatments based on subject-level coefficient estimates of the mixed regression in Equation (C4) and assuming no experimental noise (the residual is set to zero).

Treatment	Obs	Mean	Std. Dev.	Min	Max
Negative	125	0.262	0.183	0.001	0.783
Neutral	125	0.373	0.205	0.015	0.996
Positive	125	0.403	0.206	0.021	0.998

To examine the heterogeneity of the treatment effect across subjects, for each subject we compute the difference in expected allocation between the Negative and the Neutral treatments, and again the difference in expected allocation between the Positive and the Neutral treatments. This analysis allows us to compute two relative treatment effect measures per subject. We plot the distribution of these relative treatment effects in Figure C-2. Under the null that differences in behavior across treatments are noise, the distribution of relative

[33] For example, using the binomial distribution, we calculate that the stock should objectively be deemed to have a 50% chance of being the high-paying type in 29% of the rounds: This is sure to be the case in all of the first rounds, in 4 of 9 instances of the third round, and in 24/81 instances of the fifth round (and in none of the second, fourth, and sixth rounds). Adding up and dividing by six possible rounds we get (1+4/9+24/81)/6 = 0.29.

treatment effects should be symmetric about zero. The histograms in Figure 3 suggest that this roughly holds for the Positive treatment but that the same is not true for the Negative treatment.

Figure C-2: Distribution of Allocations -- Histograms
The figure reports the distribution of subjects' expected allocations in the two treatments relative to their expected allocation in the Neutral treatment. Expectations are calculated using subject-level coefficient estimates of the mixed regression in Equation (C4) and assuming no experimental noise (the residual is set to zero).

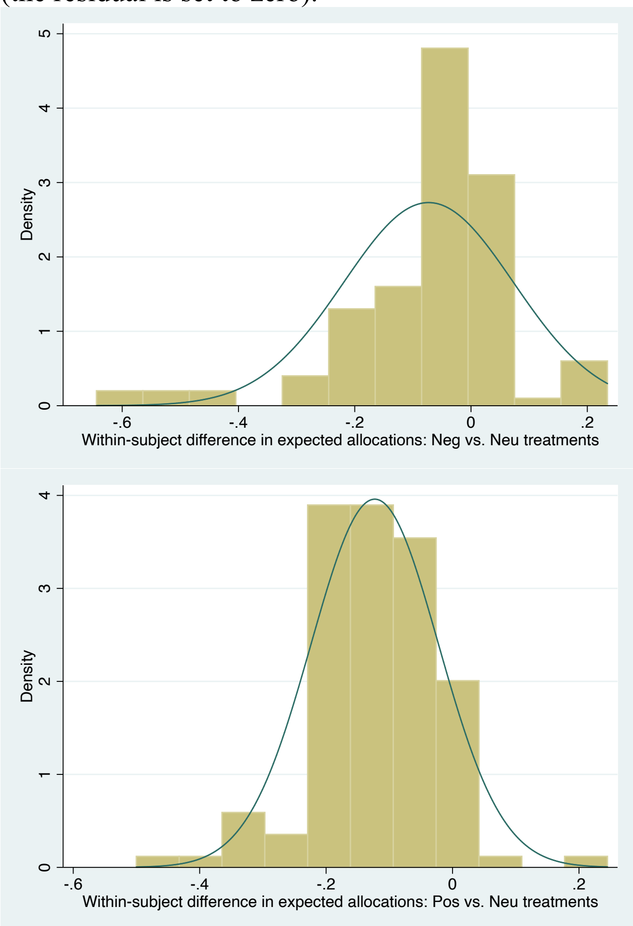

To quantify treatment-dependent heterogeneity across subjects, we estimate a 2-component finite mixture model of normal distributions to each of the histograms. The results, reported in Table C-3, suggest that the vast majority of subjects (just over 90%) are drawn from a population that expects to allocate only two more ECUs in the Positive treatment (component 1) than in the Neutral treatment, while the remaining part of the population may allocate much more (14 ECUs), but their mean allocation is not statistically distinguishable from zero (component 2). By contrast, roughly half the subjects are drawn from a population that expects to allocate 20 ECUs *less* in the Negative treatment (relative to the Neutral treatment), while the remaining expect to allocate 3 ECU less. This analysis demonstrates that the treatment effect is not just strong in aggregate but also pervasive across subjects. The finding that roughly half the subjects are highly sensitive to the Negative treatment is consistent with the survey finding by Bauer et al. (2021) that two out of three Dutch pension plan participants were in favor of RI mandates.

Table C-3: Distribution of Allocations: Finite-mixture Model Analysis

The table reports a decomposition of each histogram in Figure C-2 into a mixture of two normal distributions. Means, standard deviations, and mixture probability are estimated for the data depicted in each of the histograms. Asterisks denote the probabilities for statistical tests of differences from zero (**p*<0.1; ***p*<0.05; ****p*<0.01).

Within-subject difference in expected allocations relative to the Neutral treatment		
	Negative treatment	Positive treatment
Mean of component 1	-0.0296**	0.0211**
	(0.009)	(0.007)
Mean of component 2	-0.195***	0.12
	(0.039)	(0.144)
SD of component 1	0.0500***	0.0687***
	(0.009)	(0.009)
SD of component 2	0.274***	0.404**
	(0.027)	(0.122)
Probability of component 1	0.505	0.916***
	(0.074)	(0.052)
N	125	125

To contextualize these estimates, consider that, based on the estimates from Table C-3, 40%-50% of subjects should exhibit Type 2 tendencies in the Negative treatment and Type 1 tendencies in the Positive treatment. This translates into an additional allocation of 2 ECUs in the Positive treatment and -20 ECU in the negative treatment. The asymmetry multiplier, in this case, is ten! The remaining subjects who exhibit Type 1 tendencies in the Positive treatment will exhibit a smaller asymmetry multiplier of 1.4 that is still significantly different from 1. These two populations, both of which exhibit significant asymmetric multipliers, account for 90% of the subject pool.

Finally, to obtain further confirmation that a Negative treatment effect is pervasive among participants, we plot the expected allocation in the Negative treatment against the Positive treatment, by subject, and relative to a 45-degree line (see Figure C-3). First, the figure shows that most subjects fall under the 45-degree line, consistent with allocations in the Negative treatment being lower than in the Positive treatment, even when accounting for differences in baseline levels of allocations. Second, the asymmetry in distribution relative to the 45-degree line is observed for virtually all levels of Positive treatment allocations. Thus, the sensitivity to negative externalities is found among all subjects – those that appear to care about positive externalities and those that do not.

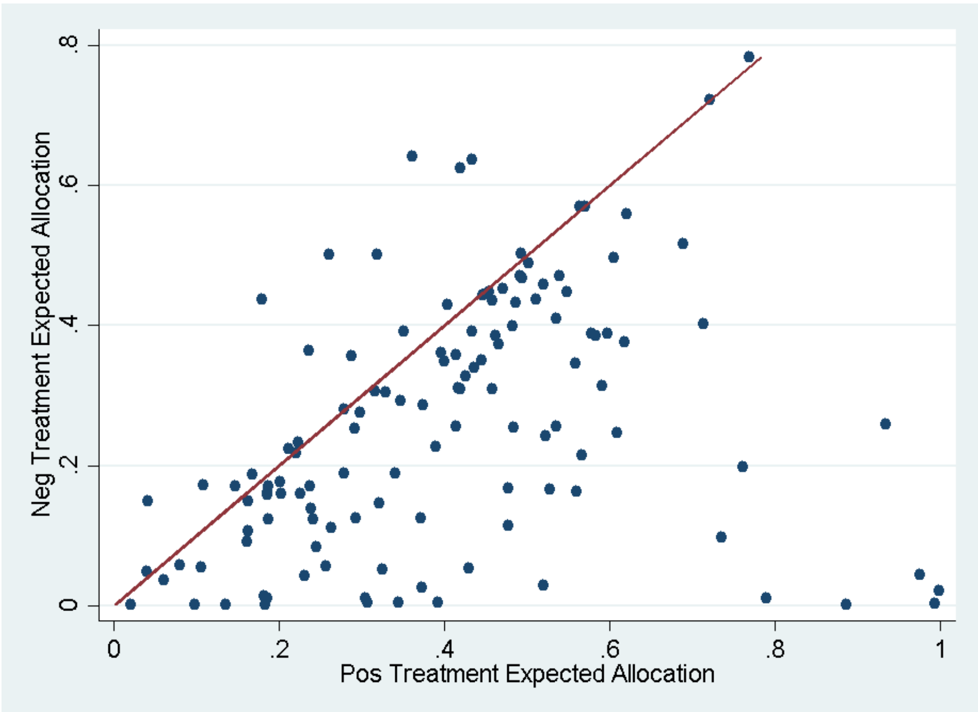

Figure C-3: Expected Allocations in Positive and Negative Treatments
The figure reports the expected allocations in the Positive and Negative treatments based on model estimates (see Table C-1). Each dot corresponds to a subject.

Appendix D: An Analysis of Learning

To further understand how considerations of externalities distort beliefs, in this analysis we examine how subjects learn in the experiment. Because probability estimates are submitted after observing realization of payoffs from the stock investment, we can investigate whether, and how, the treatments affect learning.

An extensive literature studies experimentally and theoretically how learning in various settings deviates systematically from the Bayesian framework (e.g., Tversky, 1973; Slovic and Lichtenstein, 1971; Svenson, 1981). We build on the work of Mobius et al. (2022) who nest two important deviations from Bayesian updating in a simple linear framework by transforming priors and posteriors into log odds. The nested deviations include asymmetric updating (responding differently to positive and negative signals) and conservatism (interpreting signals as less informative than they are). Specifically, we estimate the following linear regression:

$$logit(\mu_{i,t}) = \delta\ logit(\mu_{i,t-1}) + \beta_H\ I(S_{i,t} = H)\lambda_H + \beta_L\ I(S_{i,t} = L)\lambda_L + \varepsilon_{i,t} \qquad \text{(D1)}$$

where $\mu_{i,t}$ is the reported probability by subject *i* in period *t*, $logit(x) = \ln(x/(1-x))$, $\lambda_H = \ln(2)$ is the log odds of the probability that the stock is of the "High" type, $\lambda_H = -\ln(2)$ is the log odds of the probability that the stock is of the "Low" type, $I(S_{i,t} = H)$ is an indicator function for a round in which the stock doubles, and $I(S_{i,t} = L)$ is an indicator function for a round in which the stock halves. It is straightforward to check that the Bayesian case corresponds to a setting where $\delta = \beta_H = \beta_L = 1$. Any difference between β_H and β_L implies an asymmetry in updating with respect to positive versus negative information about the stock's performance. Deviation of the coefficients from one corresponds to subjects' under- or overreaction to information.

We estimate this regression for each treatment separately while clustering standard errors by subject. The regression results are presented in Table D1, panel A. Consistent with Mobius et al. (2022), we find that subjects generally tend to underweight new information when updating, as compared with the Bayesian predictions (panels A and B). In the Neutral treatment, there does not appear to be asymmetry in underweighting positive versus negative new information about the stock. This contrasts with the results in Mobius et al. (2022), potentially because, aside from financial consequences, subjects are not emotionally linked to whether the information is positive or negative.[34] Consistent with the hypothesis that the Positive and Negative treatments influence subjects' emotional connection to stock outcomes, Panel C presents evidence of asymmetry in the Positive and (especially) Negative treatments. In both cases, positive information about the stock type is discounted relative to negative news. However, the asymmetry is much more pronounced in the Negative treatment.

When we test for treatment effects, comparing the Neutral treatment to both the Positive and Negative ones, we find interesting results. First, there does not appear to be a statistically significant difference between the Neutral and Positive treatments. Second, the only significant difference that we find between the Neutral and Negative treatments is observed with respect to the response to positive information. Namely, subjects respond less to positive information in the Negative treatment relative to the Neutral one. This aligns with our observations about the role of the Negative treatment in distorting perceptions of outcomes: The Positive and Neutral treatments are nearly indistinguishable, while the Negative and Neutral treatments differ significantly.

Although the results on the asymmetry in updating in the Negative treatment are intriguing, their economic magnitude is rather small. This is consistent with the overall pattern observed in our study where belief distortion plays a secondary role when compared to the treatment effects on preferences (holding beliefs constant). To quantify this, we use the estimated coefficients to compute the posterior probability of the average subject in the Negative relative to the Neutral treatments after observing a sequence of five positive signals

[34] In Mobius et al. (2022), the information conveys personal information in that it is about a subject's IQ.

or a sequence of five negative signals. We find that after a sequence of five positive signals, subjects in the Negative treatment are predicted to report a probability of 83.4%, relative to 86.7% in the Neutral one: A 3.3% decline in reported probability. The difference after a sequence of five negative signals is only 1.8% (10.6% relative to 12.4%).

Table D1: Learning Across Treatments

The table reports the regression results corresponding to Eq. (D1), in which sequential learning is allowed to deviate from the Bayesian null for both a base-rate neglect as well as asymmetry in response to negative and positive information. The model is estimated separately for each treatment cell (Panel A); test of coefficients' deviations from the Bayesian null are included as asterisks; differences in estimated coefficients across treatments are tested (Panel B).

Panel A: Subjective Probabilities Across Treatments

	(1)	(2)	(3)
Variables	Neutral	Positive	Negative
δ	0.889**	0.905**	0.956
	(0.03)	(0.04)	(0.04)
β_H	0.674**	0.591**	0.508**
	(0.04)	(0.04)	(0.04)
β_L	0.704**	0.708**	0.672**
	(0.05)	(0.04)	(0.04)
Observations	1,159	1,169	1,125
R-squared	0.68	0.73	0.76

Panel B: Testing for Treatment Effect

(p-values)	(1)	(2)
Test	Neutral==Positive	Neutral==Negative
δ	0.6991	0.1734
β_H	0.0604	0.0006
β_L	0.9295	0.4391