NBER WORKING PAPER SERIES

THE MORAL PREFERENCES OF INVESTORS:
EXPERIMENTAL EVIDENCE

Jean-François Bonnefon
Augustin Landier
Parinitha R. Sastry
David Thesmar

Working Paper 29647
http://www.nber.org/papers/w29647

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2022

We thank participants of seminars at Harvard, HEC, MIT, Stockholm School of Economics, Toulouse School of Economics and the AFAs for their comments. Bruno Biais, Oliver Hart and Sam Hartzmark provided very useful feedback. We also thank Aidan Westley for his skillful research assistance at an earlier stage of this project. The authors are grateful to the Sustainability Initiative at MIT for funding. This study is registered in the AEA RCT Registry and the identifying number is "AEARCTR-0008764". The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Moral Preferences of Investors: Experimental Evidence
Jean-François Bonnefon, Augustin Landier, Parinitha R. Sastry, and David Thesmar
NBER Working Paper No. 29647
January 2022
JEL No. G11,G41,M14

ABSTRACT

We characterize investors' moral preferences in a parsimonious experimental setting, where we auction stocks with various ethical features. We find strong evidence that investors seek to align their investments with their social values ("value alignment"), and find no evidence of behavior driven by the social impact of investment decisions ("impact-seeking preferences"). First, the willingness to pay for a stock is a linear function of corporate externalities, and is symmetric for positive or negative externalities. Second, whether charity transfers are contingent or independent on investors buying the auctioned stock does not affect their WTP. Our results are thus compatible with a utility model where non-pecuniary benefits of firms' externalities only accrue through stock ownership, not through the actual impact of investment decisions. Finally, non-pecuniary preferences are linear and additive: willingness to pay for social externalities is proportional to the expected sum of charity transfers made by firms (even if some of these donations are negative).

Jean-François Bonnefon
Toulouse School of Economics
21, Allée de Brienne
Toulouse 31015
France
jfbonnefon@gmail.com

Augustin Landier
HEC Paris
1 rue de la libération
78351 Jouy en Josas Cedex
France
augustin.landier@tse-fr.eu

Parinitha R. Sastry
MIT Sloan
267 Beaufort Ave.
psastry@mit.edu

David Thesmar
MIT Sloan School of Management
100 Main Street, E62-632
Cambridge, MA 02142
and NBER
thesmar@mit.edu

A data appendix is available at http://www.nber.org/data-appendix/w29647

1 Introduction

This paper uses incentivized experiments to characterize investors' moral preferences. Over recent years, responsible asset management has developed considerably in size. However, the exact nature of responsible investors' preferences remains somewhat elusive. There are essentially two main views of investors' ethical preferences in the literature. One view is that responsible investors experience corporate externalities *of their portfolio companies* as a non-pecuniary dividend. This type of preferences, sometimes referred to as *value-alignment*, reflect investors' aversion for owning companies that do not have a business model in line with their own moral values. This view of investors' social preferences is the one that is most often modeled in the portfolio choice literature (see e.g. Heinkel et al. (2001) and more recently Pástor et al. (2021) and Pedersen et al. (2021)). However, a second type of social preferences might drive responsible investors : the concern for having a positive social impact (through their choices, Barber et al. (2020)). Such *impact-seeking* investors value the social consequences of their own investment decisions (the fact that these consequences exist is referred to as "additionality"). Impact preferences can be modeled by assuming that corporate externalities enter investors' utility *unconditional of the stocks they own*, with some weight. This weight reflects their degree of social responsibility (see e.g. Oehmke and Opp (2020)), Hart and Zingales (2017), Broccardo et al. (2020), Landier and Lovo (2020), Green and Roth (2021)). Philosophically, impact preferences can be associated to consequentialism (we care about consequences of our actions onto others), while value-alignment can be associated to deontological ethics (see e.g. Sandel (2007)).

Drawing a distinction between impact-seekingness and value alignment and measuring it, is particularly important for the design of financial products catering to investors' moral preferences. This is because these two sets of preferences have sensibly different implications on investment choices: for example, avoiding investing in polluting companies does not necessarily have any material impact if these companies are able to finance their projects using self-financing or other investors' capital. Impact-driven investors might value transforming "dirty companies" into less dirty ones through engagement, implying that, contrary to what happens under value-alignment, their portfolios might not necessarily be composed of companies that are more socially virtuous than average. Models of value-alignment (Heinkel et al. (2001) and more recently Pástor et al. (2021), Zerbib (2020) and Pedersen et al. (2021)) prescribe responsible products composed of companies with a better social performance than average. This leads to asset-pricing implications related to shifting demand for virtuous stocks. By contrast, models of impact investing lead to different prescriptions. For instance, in Gollier and Pouget (2014), Broccardo et al. (2020), responsible investors focus on engagement, i.e. they actively exercise control or voting rights to improve the behavior of companies. Similarly, Berk and van Binsbergen (2021) find that divestitures have only a small effect on the cost of capital of targeted companies, implying that socially conscious investors should rather invest in "dirty companies" and exercise control to change them. In Green and Roth (2021), responsible investors focus on projects that would not be financed otherwise and in Oehmke and Opp (2020), they focus on increasing the equilibrium scale of clean production, taking into account the existence of other investors. In Landier and Lovo (2020), responsible investors maximize impact on social welfare by investing in some but not all sectors and imposing pollution standards in them. These targeted sectors are not necessarily among the cleanest ones.

To illustrate how shareholder moral preferences matter for their willingness to pay for a stock, imagine a one-period setting where a company's profits per share are worth \$1. Now, assume the company is committed to spending 40% of its profits in charity donations and distributing the rest as a dividend. Non-altruistic investors would be willing to pay up to \$0.6. However, if investors value the company's prosocial behavior, the price they are willing to pay might be at a level P higher than 0.6. In this case, $P - 0.6$ measures the component of valuation by the shareholders that reflects their moral preferences. If investors are driven by value-alignment, $P - 0.6$ reflects the utility they get from holding a stock that spends 0.4 on charity donations. However, if they are impact-driven, P should be higher than 0.6 only if the donation depends crucially on them buying the stock. Indeed, if the donation is set to happen anyhow, an impact-driven investor does not feel compelled to pay more than 0.6: She will seek more efficient uses of her own generosity.

We exploit this insight by comparing investors' willingness to pay when corporate donations depends or not on them buying the stock that is auctioned. This allows us to disentangle value-alignment preferences from impact-seeking preferences. If investors only care about impact, they should not value corporate donations that would happen regardless of their investment decisions. Our key finding is that investors' valuations of corporate donations are highly significant and that they are indifferent to whether their purchase decision causes the donations or not. Thus, at least in our setting, value-alignment largely dominate and impact concerns are negligible. This result is especially striking in our setting, where prosocial impact is easily measurable (charity donations).

Telling apart value-alignment from impact-seeking preferences, as our experimental setting allows us to, is hard to do in the field. There are two main reasons for this. First, social initiatives by firms may simply be a signal of management quality or an investment in consumers and employees' loyalty, hence affecting firm value through a channel different from investors' moral preferences. Second, investors driven by value-alignment, can also, though this is not their primary objective, have an indirect impact on the behavior of companies : by avoiding investment in "dirty companies", they indirectly increase the cost of capital of such companies, hence reducing their equilibrium size and setting incentives to improve social behavior. This implies that simply observing the type of ESG policies that investors implement in the field (such as avoiding companies with poor ESG ratings) is not sufficient to disentangle value-alignment from impact-seekingness.

Let us now describe more precisely the backbone of our experiment. We elicit investors' moral preferences by auctioning several types of synthetic companies to participants: some companies are ethically neutral, some are generous (they distribute a fraction of profits to charities), and some exercise negative externalities (they reduce planned transfers to charities). The experiment is incentivized with real money (we pay on average $21, a reasonably large amount on MTurk). We first make sure participants understand the bidding mechanism and its consequences through a demanding quiz. We then find that, in our setting, participants strongly integrate social externalities into their pricing bids, even when buying a stock does not change whether the charity transfers happen (see Section 4 for a very simple framework). This is strong evidence of the existence of value-alignment preferences. The effect is quite symmetric with regard to the sign of the externality: Participants are willing to pay \$.7 more for buying a share in a firm giving one more dollar per share to charities. Symmetrically, a firm that makes profits by exercising a negative externality of \$1 on a charity is valued \$.8 less than a similar company with no externality. We find that the scaling of non-pecuniary preferences is linear: doubling the size of a social externality doubles its impact on willingness to pay (see Figure 1, more details below). To further elicit if participants care about impact, we construct a version of our experiment where bids are pivotal, i.e. that charity transfers only happen if participants buy the stock. We find that this does not change bidding behavior: Thus, we find *no* evidence of impact-seeking motives.

Finally, we provide numerous robustness checks in order to gain external validity. In particular, we find that the expression of moral preferences through bids does not crowd out more direct charitable behavior. At the end of the experiment, we offer participants an opportunity to donate to the charity, and the resulting donations are uncorrelated with the pricing of corporate ethical behavior. In addition, we find that preferences for value alignment are additive and linear. When corporate donations are uncertain, private benefits from own charitable stocks are proportional to the *expected* charity donation. When corporate action is ambiguous (externalities are positive to one charity, negative to another), private benefits are proportional to *total* charity donation.

In the following, Section 2 connects our findings to three strands of economic literature. Section 3 describes our experimental design. Section 4 develops a simple analytical framework that can be used to analyze results. Section 5 analyzes in detail our main experimental results and their economic interpretation. Section 6 studies the robustness to various features of the experimental setting. In particular, we analyze in this section the role of ambiguity and uncertainty. Section 7 concludes.

2 Literature Review

Our paper is related to two different strands of the literature. A first set of papers in behavioral economics explores how moral preferences of agents are expressed in a market context. First and foremost, a large literature on altruism documents the prevalence of "warm-glow" (or impure) altruism, i.e. that individuals derive utility from the act of donating itself rather than the impact of the donation (see for instance Ottoni-Wilhelm et al. (2017)). Our result is reminiscent of this finding except that, in our setting, individuals do not *donate* directly, but instead purchase stocks from a seller (the experimenter). Our results are closer in spirit to Elfenbein and McManus (2010), who find that, in a sample of E-Bay auctions, consumers are willing to pay higher prices for products that generate charitable donations. Bartling et al. (2014) use a lab experiment to show that in a market context consumers refrain from buying goods from firms which have negative social impact. Tasimi and Gross (2020); Tasimi and Wynn (2016); Crockett et al. (2017) also document a similar effect outside a market context, showing that people display an aversion for money earned in a manner that directly or indirectly harmed others. These moral preferences generate a price premium for socially responsible products. Leszczyc and Rothkopf (2010) also find that auctions with proceeds donated to charity lead to significantly higher selling prices, due to higher bidding from participants with charitable motives. However, some papers show that a market context tends to dampen the acuity of moral concerns. Falk and Szech (2013) documents that markets inherently erode socially responsible behavior. They use a lab experiment to measure individuals' willingness to pay for avoiding the death of the mouse, and show that this willingness to pay is lower in a market setting than in comparable non-market contexts. Sandel (2012) develops a philosophical analysis on how markets undermine moral values. We contribute to this literature by providing evidence that moral concerns strongly affect investor's willingness to pay for *financial* claims, and that investors do not take into account whether their decision to buy a stock is pivotal for the course of firms' ethical decisions.

Our results also contribute to the literature in financial economics that is concerned with socially responsible investors and their effect on stock prices and corporate policies. Using a survey, Riedl and Smeets (2017) find that "social preferences" (in the sense of reciprocity / fairness) is an important predictor of the propensity to invest in responsible financial products ; and Bolton et al. (2019) use the trail of proxy votes to infer the distribution of shareholder preferences: They find that a group of investors, including public pension funds, systematically support a more social and environment-friendly orientation of the firm. Heinkel et al. (2001) and more recently Pástor et al. (2021) and Pedersen et al. (2021) develop equilibrium models where a fraction of investors have a distaste for holding firms that are not clean. "Dirty" companies trade at a discount compared to their "clean" peers, because in equilibrium, their shareholders (i.e. those that have no moral concerns) are more concentrated in "dirty" companies. In line with this model, Hong and Kacperczyk (2009) documents that "sin stocks" exhibit positive abnormal returns. By contrast, Edmans (2011) documents that firms that treat employees relatively well have positive abnormal returns, which goes against the view that their cost of capital is lower. Margolis et al. (2007) provides a meta-analysis of the empirical literature that shows ambiguous correlations between social responsibility and financial returns. Derwall et al. (2011) finds evidence that reconciles these seemingly opposite results on returns due to the coexistence of values driven and profit-driven SRI investors. Krüger (2015) documents that stock prices react negatively to negative CSR events. Hartzmark and Sussman (2019) documents large outflows when funds are categorized as having a poor sustainability footprint. Our contribution to this literature is to isolate the stockholder preference channel for the impact of CSR on stock prices. In all these event studies, this channel is confounded with the impact of CSR news on profits (for example via employees, customers or future regulation) and so it is hard to know if CSR is priced because it enhances financial value or because shareholders value ethical behavior beyond cash-flows.

Our results are also related to concurrent and complementary work using experiments to shed light on the pricing of CSR. The key difference between these papers and ours is that we explore the distinction between value alignment and impact concerns, a crucial distinction for models and savings product design. Brodback et al. (2019) use like us an experiment to explore investor valuations of ethical behavior. Their paper focuses on whether ethical preferences are state dependent, and whether participants care more about some charities than others. They find that investors' willingness to pay for ethical behavior is lower when financial performance is poor. In their experimental setting, all participants are pivotal for the charity outcome; all participants see the same set of charities; and firms are either ethical or neutral. In contrast, our paper focuses on the valuation of ethical, neutral, *and unethical* firms, and we consider both pivotal and *nonpivotal* investors, allowing us to disentangle impact-seeking vs. value-alignment preferences. Humphrey et al. (2021) analyze behavior and learning of investors in an experiment where returns from stocks picked by the investor are positively or negatively matched by the experimenter with transfers to charities. They show that investors invest relatively less in assets when such investments have negative impact on charities. But the effect is asymmetric, in that investors do not invest more in stocks entailing positive charity transfers. Besides the absence of pivotality (our key concern), the key difference between their setting and ours is the presence of expectations formation about cash-flows and externalities. Our setting is more parsimonious: There is no learning and the preliminary quiz ensures participants have exactly the right expectations about cash-flows and donation.

3 Experimental Design

We first describe the overall structure of our experiment, and discuss details of the various conditions in the following subsections.

3.1 Overall structure

We recruit participants through Amazon's Mechanical Turk (MTurk). In our experiment, participants submit bids for shares in fictional companies that vary by how much dividend they pay to the individual, and by how much money they add to, or remove from, a fund that will be donated to charities. Participants bid on three different company types in random order. The "ethical" company gives away some shareholder profits to charity; the "unethical" company takes money away from the charity and gives it to shareholders; and the "neutral" company neither gives nor takes money from the charity wallet.

To elicit truthful valuations, we use the classic Becker-DeGroot-Marschak (BDM) bidding mechanism where participants bid in a second-price auction against a random machine (Becker et al., 1964). Under some restrictive assumptions, this ensures that participants bid their truthful valuation (see Section 4). The BDM mechanism is the following. Participants first place their bids, and then a share price is drawn from a uniform distribution after bids are submitted. Participants must then buy the share if their bid is larger than the random price. To emphasize that participants are playing against a random price, we present the random value as the result of a spinning wheel of fortune (see Appendix Figure A.2, to which we return below).

We define three variables:

1. "Selfish value" is the cash dividend directly paid to the participant by the firm. Individuals only receive a dividend from the company if they buy the share.

2. "Charity value" is the amount added to, or subtracted from, the charity wallet by the firm. In the case of an unethical company, the participant receives the company profit plus some amount subtracted from the charity wallet, and so the charity value is negative. Otherwise, it is zero (neutral company) or positive (ethical company).

3. "Excess bid" is the difference between bid and the selfish value. For pure cash-flow maximizing investors, excess bids should be zero.

As our framework in Section 4 makes clear, the key parameter of interest is the relationship (regression coefficient) between "excess bid" and "charity value". This relationship indicates to what extent individuals value the firm's behavior towards the charity, *conditional* on owning the stock. At one extreme, a one-to-one relationship between excess bids and charity values would indicate that investors completely internalize what happens to the charity. An the other extreme, excess bids of zero means that investors are only interested in cash-flows, though it may not mean that they are purely non-altruistic, as we discuss in Section 4.

3.2 Sequence of Events: Baseline Condition

We describe now the baseline condition in detail, and the variants in later sections. In the baseline, we start the experiment by asking respondents to agree to a consent form that includes a one sentence description of the experiment, a ball park estimate of payments, and the experiment's expected duration. In the first page of the interface per se, participants are given a short description of the game. They are told they will begin the experiment with a "virtual wallet" containing $2.00. Separately, $1.00 is placed in a fund that will be donated to a charity, which we refer to as the "charity wallet." Participants are then told that they will make investment decisions that affect how much money is added or subtracted from both their wallets and the charity wallet. At the end of the experiment, participants receive a base payment of $2.00 and a bonus equal to the amount in their virtual wallet. The charity receives the final content of the charity wallet.

In the baseline condition, we make it clear to the participant that both successful and failed bids lead to the same changes in the charity wallet. Thus, purely impact-driven investors should not be willing to pay more than the selfish value (the expected cash-flows). Put differently, even if they are altruistic, consequentialist investors are not expected to bid at a premium for charity-giving stocks (or at a discount for charity-taking stocks), since their actions have no consequence.

At the end of this first page, participants are asked to select the charity that will receive the content of the charity wallet. Participants choose from the following six options: the American Civil Liberties Union, the World Wildlife Fund, Food for the Poor, the American Cancer Society, Save the Children, and the Environmental Defense Fund. The charities are well-respected nationally and span a range of environmental, social, and governance issues. We provide a screenshot of the first page in Appendix Figure A.1.

Participants then proceed to the practice quiz, a key step designed to ensure that participants fully understand how the bidding game works and the consequences of their choices on their wallet and the charity wallet. Participants are first shown a detailed example of the "neutral" firm that does not modify the charity wallet. They are forced to click line-by-line through the example to ensure slow digestion of information. They do not make any decisions and are not asked any questions during the example (we provide a screenshot of this example in Appendix Figure A.2). Afterwards, participants are quizzed on both a hypothetical "ethical" and an "unethical" firm, which respectively add money to, and subtract money from, the charity wallet. Participants are given three opportunities to obtain a perfect score on the whole practice quiz. Only those who pass may continue to the main experiment. This is done to ensure that participants understand the consequences of their actions on flows of money.

The details of the practice quiz work as follows. The individual is shown, in random order, two hypothetical situations, which vary according to two dimensions. The first dimension is about the company's prosocial actions: one is ethical and the other one unethical. Specifically, the ethical company gives \$0.4 of its \$1.5 profit to charity, and the unethical company earns \$0.7 in profits and takes \$0.4 from the charity. The second dimension is about the success or failure of the hypothetical bid. In the "succeed" scenario, the random bid is set at \$0.50, while the hypothetical bid is at \$1.1, so that the company share is actually purchased (for \$0.50). In the "fail" scenario, the random bid is set at \$2, above the hypothetical bid value of \$1.1, so that the company share is not purchased. Participants see two scenarios drawn at random: The first firm is ethical with probability 1/2 (in which case the second firm is unethical), and purchased with probability 1/2 (in which case the second firm is not purchased). After presenting each hypothetical, we quiz participants on whether the firm's share is purchased, how much they would hypothetically receive in dividends, and how much the charity would hypothetically receive/lose under the given parameters. We provide a screenshot of the practice quiz for the unethical company in Appendix Figure A.3.

Once participants have fully mastered the quiz, they progress to the main experiment. In the main experiment, participants submit bids for a share in each of the three hypothetical companies: Neutral, Ethical and Unethical. Each company randomly draws a profit from $\{0.5, 0.6, 0.7, 0.8, 0.9, 1\}$. The neutral company gives the entirety of this profit to the shareholder. In contrast, the ethical company gives a random portion of this profit to the charity, drawn from $\{0.1, 0.2, 0.3, 0.4, 0.5\}$. The unethical company gives the shareholder the entirety of its profit along with money which it takes away from the charity wallet. The amount subtracted from the charity is also drawn from $\{0.1, 0.2, 0.3, 0.4, 0.5\}$. We provide a screenshot of the bid for an unethical company in Appendix Figure A.4. We randomly vary the order in which firm types (neutral, ethical and unethical) are presented in the main experiment (as we did in the quiz). At the end of the paper, we test if our results are affected by the order of presentation – and find that they are not.

After bidding on all three companies, participants are shown the amount in both their personal wallet and the charity wallet. Finally, we ask participants to answer a short survey designed to provide data on socio-demographics (education, age, gender, financial literacy).

3.3 First Batch of Experiments

We started with three pilot sessions to optimize our formulation and make sure participants could pass the quiz. The first one took place on March 20, at the BLab of MIT with 20 participants in person. The second and third pilots took place on July 3 and 17, each including 30 participants recruited through MTurk. After the second pilot, we found that quite a few participants still failed the quiz, and so we clarified the presentation of the information.

We then conducted the first batch of experiments in two waves, first on July 22, 2019, and then two weeks later on August 5th, 2019. Workers could only participate in the experiment once and were randomly assigned to different treatment groups.

In the first round of experiments (July 22nd, 2019), we recruited 300 workers and randomly allocated each worker to one of the first two conditions (*baseline with donation* and *pivotal with donation*). All participants in this batch had the option to directly donate money from their wallet to a charity of their choice at the end of the experiment. The two conditions are as follows:

1. *Baseline with donation*: This condition is the baseline described above, except that, in addition, participants are given an opportunity to directly donate to the charity from their

personal wallet. This opportunity is offered at the end of the experiment, but adverstized in the first screen. We provide a screenshot of the donation screen in Appendix Figure A.5.

2. *Pivotal with donation*: While in the baseline presented above, the charity wallet is affected whether the participant buys the stock or not, in the pivotal condition, the charity wallet is only affected if the participant buys the stock. We include this condition to test whether participants are impact-motivated or not, i.e., whether impact affects the pricing of prosocial behavior. In this condition, the quiz is adapted to reflect the effect of bid success, so that we ensure participants are fully aware of the consequences of their actions (i.e. that if the bid fails, there is no transfer to charity). Like in the first condition, donation is an option at the end, advertized at the beginning.

In the second round (August 5th, 2019), we recruit 455 workers. These workers are randomly allocated to one of three conditions (*Baseline*, *Delegation* and *BL with random charity*). This round assumes nonpivotality (like the baseline) but removes the option to donate directly (we did not observe much donation in the first wave, so decided to simplify the experiment – we will get back to this below). The three conditions of the second wave are described below in detail:

1. The pure *Baseline* condition is exactly the one described in the previous section. Since it has no donation option, it can be compared to the first condition of the first wave to investigate the effect of donations – more on this below.

2. The *Delegation* condition differs from the pure *baseline* condition in that each participant manages the wallet of another participant. They are told so explicitly at the beginning of the experiment – the personal wallet is renamed the "other person's wallet," and we verify explictly at the end of the practice quiz that people understand that their own wallet is not affected by their decision. We include this condition to check whether generosity becomes more pronounced when managing other people's money.

3. The *BL with Random Charity* differs from the *Baseline condition* in that participants do not choose which charity receives the donation. The idea here is to check whether participants only value prosocial behavior when the charity is the one they care about. We include this condition in the August 5 batch, and like other conditions in this batch, there is no option to directly donate.

Summary statistics about quiz passing rates, payments, as well as a recap on conditions, are provided in Table 1, columns 1 and 2.

3.4 Second Batch of Experiments

In the real world, social externalities are less clear-cut than in our baseline: There can be ambiguity about the behavior of companies (they may do both some social good and social harm) or perceived uncertainty about what they do. We thus decided to add new conditions to test the robustness of our results to the presence of uncertainty and ambiguity.

We ran these new experiments in two waves (June and July 2020). The COVID crisis tended to make MTurk results less reliable in the spring of 2020, so we decided to wait a little and stagger the survey for quality control (also, our quiz ensures we filter out inattentive participants). We randomize the choice of charities since, as will become clear below, the *ambiguity* condition would require to choose three different charities. Finally, for reasons that will become clear below, each condition has 6 rounds, instead of 3 in the first batch. Also, there is no donation option, like in the second wave of the first batch.

The three conditions to which participants are randomly allocated are as follows:

1. *BL with Random Charity 2* which is the baseline described above in Section 3.2, except that (1) the charity is randomly chosen and (2) there are 6 company bids instead of just 3. Companies 1,2,3 are exactly like in the baseline condition: one neutral, one ethical and one unethical in random order. After bid 3, participants start with a new charity wallet of \$1 and a new personal wallet of \$2, and do another round of 3 bids exactly like the first

ones. At the end, participants get the base payment of $2 *plus* the two personal wallets. This condition allows to control for "bidding fatigue", since uncertainty and ambiguity both require 6 bids as shown below.

2. *Uncertainty* condition. Here, the first 3 bids are identical to the first condition, they serve as some sort of warm-up session. After the third bid, participants start with a new personal wallet and one charity wallet. Then, participants face a second sequence of three bids, but now with uncertain charity value. For each company, the transfert to the charity is randomized by the toss of a coin. There are three types of companies: expected neutral (transfer is $+x$ with probability $1/2$ and $-x$ with else), expected prosocial (transfer is $+x$ with probability $1/2$ and 0 else), expected antisocial (transfer is $-x$ with probability $1/2$ and 0 else). x is drawn in $\{.1, .2, .3, .4, .5\}$. Each participant bids on each one of these company types, in random order.

3. *Ambiguity* condition. Again, the first 3 bids are identical to the first condition, as a "warm-up". After the third bid, participants start with a new personal wallet but now *two* charity wallets, say, for charities A and B. In the second sequence of 3 bids, companies transfer an amount $x_A > 0$ to A, and an amount $-x_B < 0$ to B. x_A and x_B are drawn in $\{.1, .2, .3, .4, .5\}$. There are three types of companies: ambiguous neutral ($x_A = x_B$), ambiguous prosocial ($x_A - x_B > 0$), and ambiguous antisocial ($x_A - x_B < 0$). Participants are asked to bid for each type of company in random order.

Last, for both uncertainty and ambiguity conditions, we will later introduce the notions of total and expected charity values. In the uncertainty condition, the *expected charity value* is equal to 0 for the expected neutral firm, $x/2$ for an expected prosocial firm and $-x/2$ for an expected antisocial firm. In the ambiguity condition, the *total charity value* is $x_A - x_B$ in all cases.

Summary statistics about quiz passing rates, payments, as well as a recap on conditions, are provided in Table 1, columns 3 and 4.

3.5 Descriptive Statistics

Table 1 shows summary statistics about quiz passing rates, payments and condition characteristics across the two batches and four waves of our experiment. There are 1,552 participants in total, but only 984 passed the quiz (passing rate: 63%). The passing rate is much lower in 2020 (about 50%) than in 2019 (about 80%), which is consistent with anecdotal reports that MTurkers have been providing noisier results after COVID. Among the 984 who passed the quiz, each one of our 8 conditions receives about 120 participants (see Table 2).

Bonuses range from $2.2-$2.3, and charity dividends are close to $1. In the donation round, 39 out of 255 participants choose to donate directly to the charity, donating approximately $0.54 in the baseline case and $0.64 in the pivotal case. Overall, in the first batch, participants earned about $4.2 in about 12mn on average, thus more than $20 per hour: A significantly higher pay than most MTurk jobs. In the second batch, pay was about 50% more generous.

Overall, we had a roughly even balance across demographic groups in each treatment. Table 3 presents the summary statistics on demographics by treatment condition. Panel A presents statistics on participants who failed the quiz, and Panel B presents statistics on those who passed the practice quiz and made it to the main experiment. About 60% of the participants are male, and roughly 70% report having college or graduate degrees. Table 3 also shows a difference between the first and the second batch (first 5 versus last 3 conditions), especially for college graduation, gender, financial literacy and propensity to own stocks. This difference is quite pronounced among those who failed the quiz, but less pronounced among those who made it into the experiment. This suggest that, although the composition of MTurkers was indeed affected by the COVID crisis, our quiz is successful at filtering an homogeneous group of more reliable participants.

In the remainder of the paper, we omit the delegation question from the analysis, as it does not clearly tie to our hypothesis. Although the result is intriguing and interesting (the sensitivity of bidding to charity value is the same as in the baseline), they are not directly related in our framework, to which we turn now.

4 Model

We provide here a very simple framework to understand the bidding behavior of an individual in the experiment. The model allows for both value-alignment and impact-driven preferences.

4.1 A simple model of investors' social preferences

Let u_h denote the participant's utility from holding the stock:

$$u_h = s + f(c) + \epsilon$$

where s is the dividend (selfish value) and c the charity value of the stock. ϵ is an idiosyncratic noise that varies across participants.

If the participant does not hold the stock, her utility is:

$$u_r = g(c) + \nu$$

where $g(c)$ reflects the additional utility the individual feels from money transfer c going to the charity when she does not hold the stock. We further make the light assumption that $g(0) = 0$.

The BDM mechanism relies on a second price auction against a random machine. The participant obtains the stock if her bid is above the random draw, and then pays the value of this random draw for the stock. This random bid is drawn from a uniform distribution with support $[0, \bar{p}]$. Thus, ex-ante expected utility from bidding b is given by:

$$Eu = \frac{1}{\bar{p}} \left(\int_0^b (u_h - p)dp + \int_b^{\bar{p}} (u_r)dp \right)$$

where $(u_h - p)$ is the utility when the particpant "wins the auction" against the machine, which occurs when $b > p$. Otherwise, the participant just obtains her reservation utility u_r.

Maximizing expected utility with respect to b gives us the optimal bid:

$$\underbrace{b - s}_{\text{Excess bid}} = f(c) - g(c) + (\epsilon - \nu). \tag{1}$$

In the paper, we regress $b - s$ (the "excess bid") on c (the "charity value") to identify $f(c) - g(c)$. The algebra makes clear that running this regression in the baseline case only allows us to identify the difference between the utility of corporate charity when the participant *owns* the stock, and the utility of corporate charity when the participant *does not own* the stock. For an investor whose social preferences are purely impact-driven, $f(c) = g(c)$ because holding or not the stock does not affect charity transfers, hence we expect an excess bid equal to zero for these investors. The opposite situation, of pure value alignment, is one where $g(c) = 0$: Then, the participant needs to own the stock to derive any utility from the company's prosocial behavior. But in the baseline condition, the function $g(c)$ is not identified, only the difference $f(c) - g(c)$ is.

The comparison between pivotal and non-pivotal conditions allows to pin down $g(c)$ and thus further disentangle value-alignment and impact-seeking preferences. Let us now consider the pivotal condition, where the charity does not receive/lose any money if the agent does not buy the stock. In this case, we know (by definition of g) that $u_r = g(0) = 0$ so in the pivotal condition we obtain that:

$$b - s = f(c) + (\epsilon - \nu) \tag{2}$$

Thus, in the pivotal condition – and neglecting donations which as we will see play no role – regressing the excess bid on charity value c allows to retrieve $f(c)$. Then, comparing this function with $f(c) - g(c)$ estimated in the baseline, one can infer $g(c)$. As we will see, the excess bid is as sensitive to c in the baseline as in the pivotal condition. This points to the idea that, while the participant values prosocial behavior when she owns the stock, she does not value it *at all* when she does not own it. This is 100% value alignment, in other words, $g(c) = 0$.

4.2 Empirical Specifications

We use i to index individuals and j to index bidding round. We denote *excess bids* as $e_{ij} = b_{ij} - s_{ij}$. To simplify exposition, we assume linearity: $f(c) = \alpha c$ and $g(c) = \beta c$. The functional forms are identified, but they will turn out to be linear in the data (see Figure 1 and analysis below). In this case, Equation (1) yields the empirical relationship:

$$\underbrace{e_{ij}}_{ExcessBid} = (\alpha - \beta) \underbrace{c_{ij}}_{CharityValue} + \varepsilon_{ij} \tag{3}$$

As noted above, the coefficient measures the spread between pro-social preferences when owning (α) and not owning (β) the stock.

In the pivotal condition, where charity transfers happen only conditional on the investor winning the auction, the above analysis yields the following empirical specification:

$$e_{ij} = \alpha c_{ij} + \varepsilon_{ij}. \tag{4}$$

These equations make clear how identification works. In line with Leszczyc and Rothkopf (2010), we can infer different charitable motives based on our estimates of α and β. A result of $\alpha = \beta = 0$ indicates a purely selfish model in which bidders do not internalize what happens to the charity at all. If $\alpha = \beta > 0$, bidders obtain utility from money going to charity, but are indifferent about the source of the donation – they only care about impact. Alternatively, if $\beta > \alpha = 0$, bidders are 100% value aligned: They only care about the firm's prosocial behavior when they own the stock.

5 Main Results

In this section, we first explore the baseline treatment, where the charity value of the stock does not depend on whether the participants owns it or not. We then compare the baseline treatment with the pivotal treatment, which allows to separately identify α and β discussed in Section 4, the sensitivity to charity value when they stock is owned or not.

5.1 Baseline: Evidence of value alignment

We first establish linearity by putting together several baseline-like conditions. We pool together: Baseline, BL with donation, BL with Random Charity and BL with Random Charity 2 (we only take the first 3 bids). The precise set of conditions chosen does not affect the overall picture since, as we will see later, charity randomization or donation have no effect on the sensitivity of bidding to charity value. But we thought it was best to start with the largest possible baseline. Together, these baseline conditions feature 481 participants, or 1,443 rounds of bidding.

We show in Figure 1 the average excess bid (bid minus selfish value) as a function of charity value across all four conditions. The relationship is strikingly linear. This relationship is robust to adding more conditions or removing some of them. The blue line if a fitted linear regression: Its slope is about .8 and is highly significant ($t = 20$, after clustering at the participant level).

We report formal regression results of Equation 3 in Table 4. Again, here, we focus on the same 4 quasi-baseline conditions. The regression coefficient is expected to be zero if $\alpha = \beta$, i.e. when participant are purely driven by impact. One particular case of this is non-altruistic participants who are such that $\alpha = \beta = 0$. Consistent with Figure 1, we can reject this hypothesis very strongly: participants incorporate the charity externality into their bidding behavior, and much more more so if they end up owning the stock ($\alpha > \beta$). Column (1) shows that a \$1 increase in charity dividends translates to a \$0.78 increase in bids above selfish values. In the context of the model, this gives us an estimate of $\alpha - \beta = 0.78$, which represents the additional utility shareholders feel from money going to the charity if they own the stock. Columns (2)-(3) confirm that the slope is similar in the first and second batch of our experiments, which is reassuring since COVID is said to have affected the performance of MTurkers in online experiments. It is likely that our quiz allowed to control quality post COVID pretty effectively.

Columns (4) and (5) test linearity. Column (5) of Table 4 explores the possibility that charitable giving may have asymmetric effects. We can test for asymmetry because each participant faces a prosocial company ($c > 0$) and an antisocial one ($c < 0$). One hypothesis could be that participants discount prices of companies that take money from charities more than they value charitable donations (a form of prospect theory loss aversion). There is some evidence of loss aversion, but it is not significant (p value of equality of positive and negative effects of .247). Our results are at odds with Humphrey et al. (2021), whose setting is more complex than ours (it has learnings and multiple bidding opportunities per firm).

5.2 Pivotal condition: Evidence of no impact-seekingness

Our baseline condition only identifies $\alpha - \beta$. The result that $\alpha - \beta > 0$ indicates that value alignment is present in our participants' social preferences ($\alpha > 0$), but does not prove that impact-seekingness is absent: It could still be that $\beta > 0$, i.e. that participants care about the company's prosocial behavior even when they do not own the stock.

As discussed in Section 4, in order to separately identify α and β, one can compare the pivotal and the non-pivotal conditions. To make the comparison as tight as possible, we focus here on two conditions that are exactly comparable except for pivotality: *Baseline with donation*, and *Pivotal with donation*. Both conditions were run at exactly the same time, have exactly the same structure except that in one of them the company transfers to the charity only in the event of buying.

We report the result in Table 5, we regress excess bids on charity value in both conditions separately, and compare the two. Our baseline estimates in Column (1), assuming non-pivotality, give us $\alpha - \beta \approx 0.8$, which is nearly identical to the slope coefficient in Table 4 (which uses a slightly more extensive definition of the baseline by including random charities and no-donation baseline). In column (2), we focus on the pivotal condition. There, the coefficient $\alpha \approx .9$, so, indeed, this leads to $\beta \approx .1 > 0$, but this difference is not significant (p value .345). Thus, participants only seem marginally interested in the prosocial behavior of firms even if they do not own them, if at all. We report in Figure 2 a graphical presentation of this finding.

Results from Table 5 lend support to the hypothesis that people are near 100% value-alignment driven: They only get utility from money given to charities by stocks they hold.

6 Robustness

This section explores how our results are affected by various changes in our experimental setting. In the real world, firms prosocial actions are rarely as pure as in our experiment. They may be offset by other anti-social actions (as in the case of carbon offsets), they may be uncertain. Prosocial actions of firms may not be fully aligned with investors' preferences. Finally, shareholders have the opportunity to be generous by themselves, which may crowd out the WTP for virtuous companies. We explore these deviations in what follows.

6.1 Substitution between personal and corporate donation

In the real world, individual have other ways of expressing their altruistic preferences than buying stocks. A possibility is that participants may substitute corporate charity with direct donation, as in the famous example given by Friedman (1970). In Table 6, we ask if the propensity to directly donate responds to the total charity value, which varies randomly across participants. The RHS variable, *Total charity value*, is the sum of all transferts to charity in all three bids experienced by the participants (thus, the sum of the donations of the ethical and and unethical firms, since the neutral firm does not donate). In this Table, we of course focus only on the conditions where donation is possible, i.e. *baseline with donation* and *pivotal with donation* (255 participants). In column 1, we regress a dummy equal to 1 if donation is positive (which happens in 39 cases only) on total charity value. In column 2, we use as LHS variable the total amount donated conditional on positive donation. In both cases, the correlation with total firm prosocial behavior is insignificant. Table 6 suggests that the WTP for virtuous stocks is *not* crowded out by the possibility to donate directly.

Table 7, columns 1 and 2, is another exploration of this potential substitution. Columns 1 and 2 report the estimates of Equation 3 under two conditions: *baseline* (with chosen charity and no donation allowed, 114 participants) and *baseline with donation* (with chosen charity but donation allowed, 131 participants). All other features of these two conditions are similar, except that the baseline with donation was run on July 22, and the no donation condition was run on Aug 5 of 2019. If personal and corporate charity giving are to some extent substitutes, we would expect excess bidding to be less sensitive to charity value in the donation condition – since participants can "make up" for their lukewarmness in bidding for stocks by donating afterwards. Looking at columns 1 and 2, we find essentially no difference (p-value=.143). This confirms results from Table 6 that participants view the decision to donate and the decision to "reward" pro-social companies as unrelated to one another.

6.2 Picking the charity

In the real world, investors do not get to choose the specific type of prosocial action that companies may implement. In Table 7, columns 3 and 4, we evaluate the effect of randomizing the charity on the valuation of charity value. In column 3, we focus on the *baseline* condition with no donation and charity choice (114 participants, 342 rounds of bidding). In column 4, we focus on the *BL with random charity* condition (121 participants, 363 rounds). The two conditions belong to the second wave of batch 1, and only differ in that the second one picks the charity at random. We find that excess bids are 0.31 higher (not lower) per dollar when individuals do not choose the charity, and that this difference is significant at the 1% level. Thus, we can conclude that the treatment effect is robust to removing the choice of charity. Why the result is actually stronger is not entirely obvious to us, and would warrant further investigation.

6.3 Uncertainty and ambiguity about the pro-social outcome

So far we have considered only firm donations that are certain (probability 1 or 0) and that have a unilateral dimension (firms do not exercise simultaneously positive and negative externalities). In this sub-section, we relax these restrictions and explore how the valuation of corporate externalities extends to uncertain and ambiguous situations. This is based on the second batch of experiments where we include an ambiguity condition (donations to two different charities of opposite signs) and an uncertainty condition (rolling a dice after bidding stage to determine the charity transfer). As we will see, private benefits of corporate donations are additive and linear; They aggregate nicely.

We report results in Table 8. Columns 1-2 focus on uncertainty. Column 1 uses the *BL with random charity 2* condition, while column 2 uses the *uncertainty* condition. Both conditions are identical, except for the presence of uncertainty. In column 1, we use the plain charity value as the RHS variable. In column 2, we use the *expected charity value* as the RHS variable, which is the mathematical expectation of charity transfer at the time of bidding. In both cases, we only use rounds 4,5,6 since there is no uncertainty in the first 3 rounds (which serve as a warm-up session).

Comparing columns 1 and 2, we find that the pricing of expected uncertain charity value is smaller than the equivalent certain amount, but the difference is insignificant (p=.521). Hence, there is a small discount applied to uncertain corporate prosocial behavior, but it is statistically insignificant. Note that the coefficient are a notch smaller than in the baseline regressions of Table 4, even in the baseline condition of column 1: This may be because we are focusing here on rounds 4,5,6 while results from Table 4 are on the first 3 rounds.

Similarly, columns 3 and 4 of Table 8 show that bidders do not penalize ambiguous behavior, but perform the addition of pains and gains mathematically. In column 3, we use the last 3 rounds of the *BL with random charity 2* condition; In column 4, the last 3 rounds of the *ambiguity* condition. In column 3, we use as the RHS the charity value. In column 4, we use the *total* charity value across the two charities A and B. We find that the coefficient is similar, though a little smaller. Again, the discount for ambiguous behavior may exist but it is small and insignificant.

In Appendix Table A.1, we replicate this analysis but split down *expected charity value* into $\frac{1}{2}$ charity value in case of tail and $\frac{1}{2}$ charity value in case of head, and *total charity value* into positive and negative donations. We cannot reject the hypothesis that private benefits of owning the stock are linear. This table gives slightly noisier results but confirms that value alignment preferences are additive.

6.4 Robustness to details of the experimental setting

Furthermore, in light of the extensive results on priming and framing in the behavioral economics literature, we consider how the order in which the companies are presented may affect bidding behavior. For example, it could be the case that individuals who see the unethical company first may bid lower throughout the experiment, or perhaps that individuals who see the ethical company first are induced to be consistently more generous. Besides the experiment itslef, in the practice quiz, one may be concerned that individuals who are first exposed to a situation where they managed to *buy* the share could prime them towards bidding higher to buy shares in the subsequent experiment.

To test for the possibility of priming, we varied the order in which companies are presented to the individual in both the practice quiz and the main experiment. We then run the specification in Equation 1 for each subsample, and test for equality of coefficients. Figure 3 shows that the coefficients of the regressions are very similar across practice quiz scenarios, ranging from 0.78 to 0.89. Our F-test of joint significance does not allow us to reject the hypothesis that the coefficients are all the same. In Figure 1, we see that the range of coefficients for the main experiment is a bit wider, ranging from 0.68 in the case where the ethical company is seen first, to 1.01 in the case where the neutral company is first. However, we still cannot reject the hypothesis that all the coefficients are equal. We repeat this regression within each treatment group and also find that we cannot reject the equality across coefficients (not shown here). Taken together, the robustness checks imply that our treatment effects cannot be explained by priming.

7 Conclusion

To what extent do shareholders value ethical behavior? In this paper, we develop a theoretical framework and an experimental design to investigate this question. We present a lab experiment that allows participants to submit bids for companies that vary by how much money they add or subtract from a fund that will be donated to charity. This design allows us to isolate the impact of a firm's externalities on investor bids, a feature that is difficult to achieve outside an experimental setting. We find strong evidence that individuals incorporate the charity externality in their bids and that this relationship is symmetric and almost linear. This result persists across treatment conditions, including when shareholders are pivotal for what happens to the charity, which suggests that their ethical preferences should be understood as a form of value alignment rather than impact seekingness.

References

Barber, Brad, Adair Morse, and Ayako Asuda, "Impact Investing," *Journal of Financial Economics*, 2020.

Bartling, Björn, Roberto A Weber, and Lan Yao, "Do markets erode social responsibility?," *The Quarterly Journal of Economics*, 2014, *130* (1), 219–266.

Becker, Gordon M., Morris M. Degroot, and Jason Marshak, "Measuring utility by a single-response sequential method," *Behavioral Science*, 1964, *9* (3).

Berk, Jonathan and Jules H van Binsbergen, "The impact of impact investing," *Available at SSRN 3909166*, 2021.

Bolton, Patrick, Tao Li, Enrichetta Ravina, and Howard L Rosenthal, "Investor ideology," Technical Report, National Bureau of Economic Research 2019.

Broccardo, Eleonora, Oliver D Hart, and Luigi Zingales, "Exit vs. voice," Technical Report, National Bureau of Economic Research 2020.

Brodback, Daniel, Nadja Guenster, and Sébastien Pouget, "The Valuation of Corporate Social Responsiblity: An Experiment," 2019.

Crockett, Molly J, Jenifer Z Siegel, Zeb Kurth-Nelson, Peter Dayan, and Raymond J Dolan, "Moral transgressions corrupt neural representations of value," *Nature Neuroscience*, 2017, *20* (6), 879–885.

Derwall, Jeroen, Kees Koedijk, and Jenke Ter Horst, "A tale of values-driven and profit-seeking social investors," *Journal of Banking & Finance*, 2011, *35* (8), 2137–2147.

Edmans, Alex, "Does the stock market fully value intangibles? Employee satisfaction and equity prices," *Journal of Financial economics*, 2011, *101* (3), 621–640.

Elfenbein, Daniel W and Brian McManus, "A greater price for a greater good? Evidence that consumers pay more for charity-linked products," *American Economic Journal: Economic Policy*, 2010, *2* (2), 28–60.

Falk, Armin and Nora Szech, "Morals and markets," *science*, 2013, *340* (6133), 707–711.

Friedman, Milton, "The Social Responsibility of Business is to Increase its Profit," *The New York Times*, 1970.

Gollier, Christian and Sébastien Pouget, "The" washing machine": Investment strategies and corporate behavior with socially responsible investors," 2014.

Green, Daniel and Benjamin Roth, "The allocation of socially responsible capital," *Available at SSRN 3737772*, 2021.

Hart, Oliver and Luigi Zingales, "Companies Should Maximize Shareholder Welfare Not Market Value," *Journal of Law, Finance, and Accounting*, 2017, *2* (2), 247–274.

Hartzmark, Samuel M and Abigail B Sussman, "Do investors value sustainability? A natural experiment examining ranking and fund flows," *Journal of Finance*, 2019.

Heinkel, Robert, Alan Kraus, and Josef Zechner, "The effect of green investment on corporate behavior," *Journal of financial and quantitative analysis*, 2001, *36* (4), 431–449.

Hong, Harrison and Marcin Kacperczyk, "The price of sin: The effects of social norms on markets," *Journal of Financial Economics*, 2009, *93* (1), 15–36.

Humphrey, Jacquelyn, Shimon Kogan, Jacob Sagi, and Laura Starks, "The asymmetry in responsible investing preferences," Technical Report, National Bureau of Economic Research 2021.

Krüger, Philipp, "Corporate goodness and shareholder wealth," *Journal of financial economics*, 2015, *115* (2), 304–329.

Landier, Augustin and Stefano Lovo, "ESG Investing: How to Optimize Impact?," *HEC Paris Research Paper No. FIN-2020-1363*, 2020.

Leszczyc, Peter T.L. Popkowski and Michael H. Rothkopf, "Charitable Motives and Bidding in Charity Auctions," *Management Science*, 2010, *56* (3), 399–413.

Margolis, Joshua D, Hillary Anger Elfenbein, and James P Walsh, "Does it pay to be good? A meta-analysis and redirection of research on the relationship between corporate social and financial performance," *Ann Arbor*, 2007, *1001*, 48109–1234.

Oehmke, Martin and Marcus M Opp, "A theory of socially responsible investment," 2020.

Ottoni-Wilhelm, Mark, Lise Vesterlund, and Huan Xie, "Why Do People Give? Testing Pure and Impure Altruism," *American economic review*, 2017, *107* (11), 3617–3633.

Pástor, Luboš, Robert F Stambaugh, and Lucian A Taylor, "Sustainable investing in equilibrium," *Journal of Financial Economics*, 2021, *142* (2), 550–571.

Pedersen, Lasse Heje, Shaun Fitzgibbons, and Lukasz Pomorski, "Responsible investing: The ESG-efficient frontier," *Journal of Financial Economics*, 2021, *142* (2), 572–597.

Riedl, Arno and Paul Smeets, "Why do investors hold socially responsible mutual funds?," *The Journal of Finance*, 2017, *72* (6), 2505–2550.

Sandel, Michael J, *Justice: A reader*, Oxford University Press on Demand, 2007.

_ , *What money can't buy: the moral limits of markets*, Macmillan, 2012.

Tasimi, Arber and James J Gross, "The dilemma of dirty money," *Current irections in Psychological Science*, 2020.

_ **and Karen Wynn**, "Costly rejection of wrongdoers by infants and children," *Cognition*, 2016, *151*, 76–79.

Zerbib, Olivier David, "A Sustainable Capital Asset Pricing Model (S-CAPM): Evidence from green investing and sin stock exclusion," *Available at SSRN 3455090*, 2020.

Figures

Figure 1: Linearity of the Relationship between Excess Bids and Charity Values

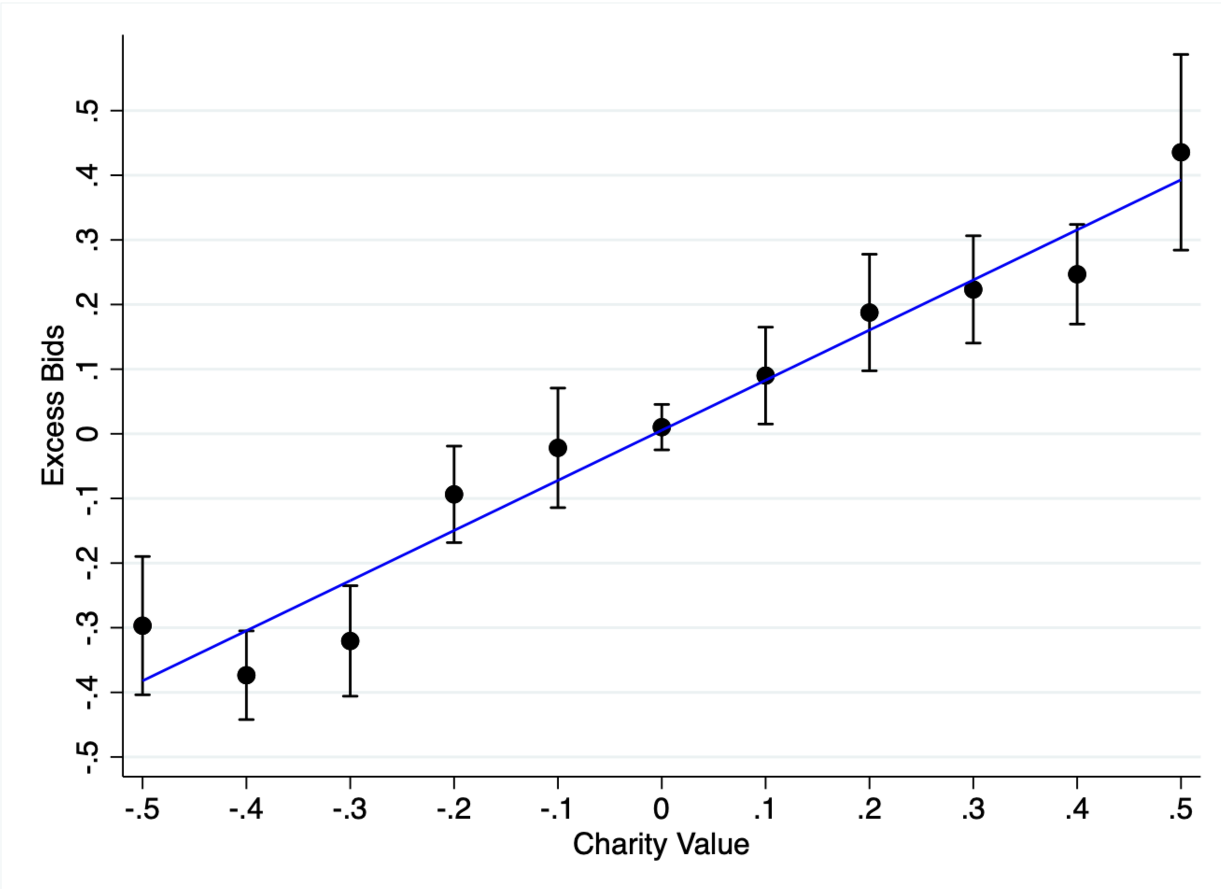

Note: We report here the average excess bids for each level of charity value between -.5 and +.5, along with 90% confidence intervals. We pool together all baseline-related conditions: Baseline. BL with donation, BL with random charity, BL with random charity 2 (only the first 3 rounds). The first three conditions were obtained in the summer of 2019; the last one in the summer of 2020. The blue line if a fitted univariate regression, which has a slope of .8.

Figure 2: Excess Bids and Charity Values: Pivotal v Baseline

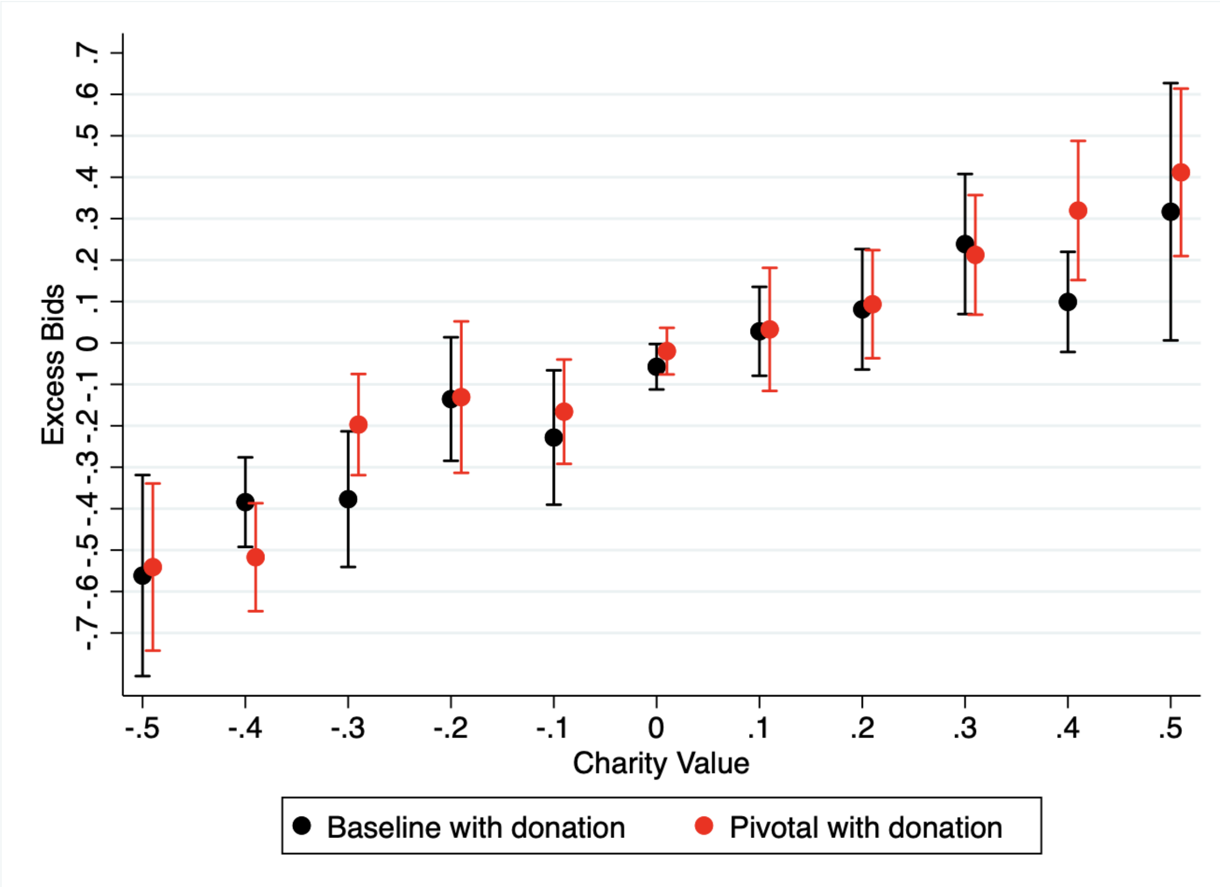

Note: We report here the average excess bids for each level of charity value between -.5 and +.5, along with 90% confidence intervals. The black chart is executed on the *baseline with donation* condition and the red chart on the *pivotal with donation* condition. Both conditions were run during the first wave of the first batch (Jul 22, 2019).

Figure 3: Robustness 2: Practice Quiz

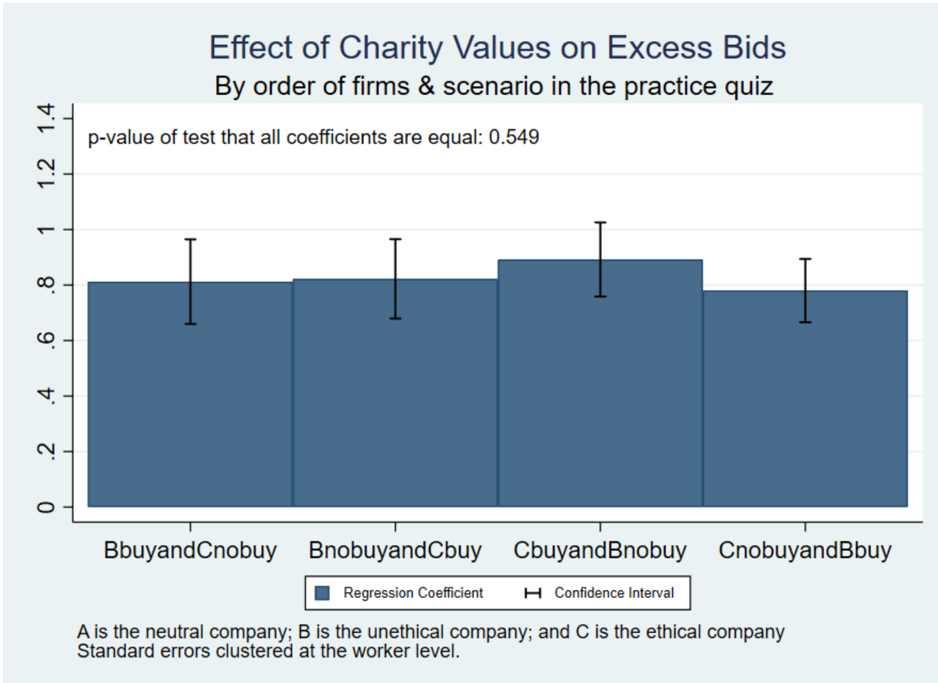

Figure 4: Robustness 3: Main experiment

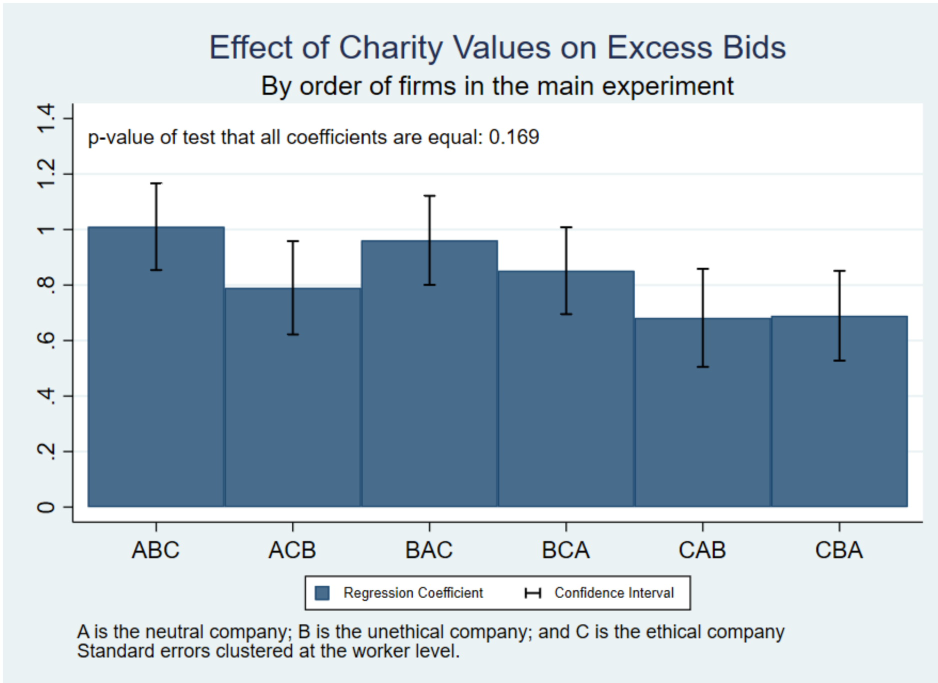

Tables

Table 1: Description of Experiment Conditions and Overall Outcomes

	First Batch		Second Batch	
	jul2019	aug2019	jun2020	jul2020
Participants (count)	300	457	448	347
Practice Quiz Statistics:				
Passed (count)	255	348	205	181
passing rate (%)	.85	.761	.458	.522
Passed in the first attempt (%)	.325	.319	.268	.326
Passed in the first two attempts (%)	.788	.813	.824	.823
Average duration (mins)	11.7	12.4	14.4	16
Base payment:				
Base Payment	2	2	2	2
Average amount in the participant's wallet:				
Final amount in participant wallet	2.3	2.3	4.47	4.42
After the first 3 companies	.	.	2.23	2.23
After the last 3 companies	.	.	2.24	2.19
Average amount in the charity wallet:				
First charity's wallet	1.05	1.01	.998	.997
Second charity's wallet	.	.	1.02	1.02
Third charity's wallet	.	.	.99	1.01

Note: This Table reports quiz performance and average pay for each batch and wave of our experiment. 1,552 persons participated in our experiment (first line), but only 989 passed the quiz (second line).

Table 2: List of Conditions

	Number of participants	Donation	Charity randomization	Number of bids
Panel A: First Batch (summer 2019)				
Baseline with donation	131	Y	N	3
Pivotal with donation	124	Y	N	3
Baseline	114	N	N	3
Delegation	112	N	N	3
BL with Random Charity	121	N	Y	3
Panel B: Second Batch (summer 2020)				
BL with Random Charity 2	124	N	Y	6
Ambiguity	113	N	Y	6
Uncertainty	145	N	Y	6

Note: This Table describes the 8 experimental conditions that we ran. We do not discuss the delegation condition in the paper to stay focused on message, but we analyze the other 7.

Table 3: Socio-demographics by condition

Condition	18-35	35-55	55^+	College	Male	Fin. Lit.	Stocks
Panel A: Did not pass							
Baseline with Donation	.64	.29	.05	.64	.76	.58	.23
Pivotal with Donation	.89	.07	.03	.64	.53	.46	.28
Baseline	.7	.23	.05	.79	.55	.38	.44
Delegation	.59	.31	.09	.77	.61	.47	.4
BL with Random Charity	.64	.25	.09	.67	.64	.35	.32
BL with Random Charity 2	.63	.31	.05	.92	.77	.18	.62
Ambiguity	.61	.32	.06	.91	.74	.14	.63
Uncertainty	.61	.32	.06	.9	.72	.16	.59
Panel B: Passed							
Baseline with Donation	.61	.31	.07	.63	.57	.77	.25
Pivotal with Donation	.62	.34	.03	.5	.62	.8	.19
Baseline	.58	.28	.12	.68	.6	.76	.28
Delegation	.6	.35	.03	.61	.55	.73	.27
BL with Random Charity	.48	.36	.14	.72	.58	.74	.29
BL with Random Charity 2	.6	.31	.07	.74	.66	.84	.31
Ambiguity	.55	.37	.07	.8	.51	.68	.28
Uncertainty	.54	.36	.08	.7	.57	.66	.29

Note: This Table describes the participants of the 8 experimental conditions that we ran. We do not discuss the delegation condition in the paper to stay focused on message, but we analyze the other 7.

Table 4: The sensitivity of Excess Bids to Charity Values: Baseline

	LHS: Excess bid				
	(1)	(2)	(3)	(4)	(5)
Charity Value	.78***	.81***	.7***	.77***	
	(20)	(18)	(9)	(20)	
$\frac{1}{2}$ Charity Value^2				-.22	
				(-.74)	
(Charity Value)$^-$					.84***
					(13)
(Charity Value)$^+$					.71***
					(9.8)
Constant	.0054	-.028*	.1***	.012	.017
	(.36)	(-1.7)	(3.1)	(.69)	(.96)
Observations	1,470	1,098	372	1,470	1,470
R^2	0.17	0.19	0.14	0.17	0.17
Batch	Both	First	Second	Both	Both
P value asymmetry					.247

Note: The LHS is the excess bid, $b - s$, the difference between the bid and the dividend (Selfish value). The RHS is c, transfered by the charity to the company (Charity value). We pool together all baseline-related conditions: Baseline. BL with donation, BL with random charity, BL with random charity 2 (only the first 3 rounds). The first three conditions were obtained in the summer of 2019; the last one in the summer of 2020. Standard errors in parentheses, clustered at the worker level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: The sensitivity of Excess Bids to Pivotality

	(1) Baseline	(2) Pivotal
Charity Value	.8***	.89***
	(11)	(12)
Constant	-.07***	-.036
	(-2.7)	(-1.4)
Observations	393	372
R^2	0.21	0.28
Batch	First	First
P value of equality	.345	

Note: The LHS is the excess bid, $b - s$, the difference between the bid and the dividend (Selfish value). The RHS is c, transfered by the charity to the company (Charity value). Column 1 uses the "baseline with donation" condition, and column 2 the "pivotal with donation" condition. Standard errors in parentheses, clustered at the worker level. * p < 0.10, ** p < 0.05, *** p < 0.01

Table 6: Effect of corporate externality on propensity to donate

	(1) Y/N	(2) Amount
Total Charity Value	.097	.19
	(.76)	(.53)
Constant	1.2***	.58***
	(50)	(7.2)
Observations	255	39
R^2	0.00	0.01
Batch	First	First

Note: The data used here corresponds to the two conditions with the option to donate: Baseline with donation and pivotal with donation, which corresponds to 255 participants. In both regressions, the RHS is the sum of charity values across all three bids. In column 1, the LHS variable is a dummy equal to 1 if the participant donates at the end. In column 2, the LHS is the amount given, conditional on giving (thus only for 39 participants out of 255). Standard errors in parentheses, clustered at the worker level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Sensitivity to option to donate and charity randomization

	Effect personal donation		Effect of random charity	
	(1)	(2)	(3)	(4)
	Baseline	Charity w. donation	Baseline	BL with random charity
Charity Value	.64***	.8***	.64***	.95***
	(8.5)	(11)	(8.5)	(12)
Constant	.0044	-.07***	.0044	-.014
	(.16)	(-2.7)	(.16)	(-.41)
Observations	342	393	342	363
R^2	0.14	0.21	0.14	0.22
Batch	First	First	First	First
P value of equality	.143		.004	

Note: In all regressions, the LHS is excess bid, and the RHS is charity value. Columns 1-2 ask if private donation is a substitute for corporate donation. Column 1 uses the *baseline* condition (with charity choice and no donation; 114 participants, 342 rounds) and column 2 uses the *baseline with donation* condition (with charity choice; 131 participants, 393 rounds). Column 3-4 evaluate the effect of randomizing the charity. Column 3 uses the *baseline* condition (with charity choice and no donation; 112 participants, 336 rounds) and column 4 uses the *BL with random charity* condition (with no donation; 121 participants, 363 rounds). We report in the last line the p-value of a test of equal regression coefficients for each comparison. Standard errors in parentheses, clustered at the worker level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Effect of uncertainty and ambiguity of corporate externalities

	(1) Baseline	(2) Uncertainty	(3) Baseline	(4) Ambiguity
Expected Charity Value	.6***	.51***		
	(7.8)	(4.3)		
Total Charity Value			.6***	.45***
			(7.8)	(3.5)
Constant	.07**	.16***	.07**	.07**
	(2.5)	(4.8)	(2.5)	(2)
Observations	372	435	372	339
R^2	0.13	0.02	0.13	0.03
Batch	Second	Second	Second	Second
P value of equality	.521		.329	

Note: In all regressions, the LHS is excess bid, and the RHS is the relevant version of charity value. In all regressions, we only use the last 3 rounds of bidding of each condition (since the first 3 rounds are just the baseline). Columns 1-2 ask if charity uncertainty affects the propensity to bid. Column 1 uses the *BL with random charity 2* condition (124 participants) and column 2 uses the *uncertainty* condition (145 participants). The RHS variable is *expected charity value*, i.e. the mathematical expectation of the charity transfer computed every round (in column 1, the baseline, there is no uncertainty so expected charity value is charity value). Column 3-4 evaluate the effect of ambiguous prosocial behavior. Column 3 uses the *BL with random charity 2* condition (124 participants) and column 4 uses the *ambiguity* condition (113 participants). The RHS variable is *Total charity value*, i.e. the sum of charity transfers computed every round (since in this condition the firm donates to one charity and takes from another one). In column 3, the baseline, there is no ambiguity so total charity value is just the charity value. We report in the last line the p-value of a test of equal regression coefficients for each comparison. Standard errors in parentheses, clustered at the worker level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$