Journal of Financial Economics 163 (2025) 103955

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/finec

The moral preferences of investors: Experimental evidence☆

Jean-François Bonnefon [a], Augustin Landier [b,*], Parinitha Sastry [c], David Thesmar [d,e,f]

[a] TSE and IAST, France
[b] HEC Paris, France
[c] Columbia Business School, United States
[d] MIT Sloan School of Management, United States
[e] NBER, United States
[f] CPER, United Kingdom

ARTICLE INFO

Dataset link: Moral Preferences of Investors : Experimental Evidence -- Replication Package (Original data)

JEL classification:
G11
G41
D62
D64

Keywords:
Portfolio choice
Externalities
Behavioral finance
ESG investing

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 (WTP) for a stock is an increasing and quasi-linear function of corporate externalities. Second, this WTP does not change when corporate externalities are made contingent on investors buying the auctioned stock. Our results are thus compatible with a utility-maximization model where non-pecuniary benefits of firms' externalities only accrue through stock ownership, not through the actual impact of investment decisions. Finally, the ability to directly contribute to the externality (by donating) does not reduce the willingness to pay for virtuous stocks.

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 preference, sometimes referred to as *value-alignment*, reflects 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 preference 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 often referred to as "additionality"). Impact preferences can be modeled by assuming that corporate externalities enter investors' utilities *unconditional of the stocks they own*, with some weight. This weight reflects their degree of pro-social preferences (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 with consequentialism (we only care about the consequences of our actions onto others, see e.g. Singer (2016)), while value-alignment can be associated with deontological ethics (we adhere to a rule independently of the

☆ We thank Dimitris Papanikolaou who was the editor for this article. We also thank participants of seminars at Harvard, HEC Paris, 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. Aidan Westley provided 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".

* Corresponding author.
E-mail addresses: jean-francois.bonnefon@tse-fr.eu (J.-F. Bonnefon), landier@hec.fr (A. Landier), prs2107@columbia.edu (P. Sastry), thesmar@mit.edu (D. Thesmar).

https://doi.org/10.1016/j.jfineco.2024.103955
Received 16 September 2022; Received in revised form 8 October 2024; Accepted 10 October 2024
Available online 1 November 2024

consequences of our own actions, see e.g. Sandel (2007)). In this paper, we use the terms impact-seeking/consequentialist and value alignment/deontological interchangeably.

Distinguishing between impact-seeking behavior and value alignment, as well as 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 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 themselves or with other investors' capital. Impact-driven investors 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) and 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 on them. These targeted sectors are not necessarily among the cleanest ones.

To illustrate how shareholders' moral preferences matter for their willingness to pay for a stock, let us consider the following thought experiment. Imagine a one-period setting where a company's profit per share is worth \$1. Now, assume the company is committed to spending 40% of its profit on 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 on them buying the stock. Indeed, if the donation is set to happen anyway, an impact-driven investor does not feel compelled to pay more than 0.6: She will seek to be more efficiently generous.

We exploit this insight by comparing investors' willingness to pay when corporate donations either do or do not depend 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 or not their purchase decision causes the corporate donation. Thus, at least in our setting, value-alignment largely dominates and impact concerns are negligible. This result is especially striking since in our setting 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 for disentangling value-alignment from impact-seeking behavior.

Let us now describe more precisely the backbone of our online 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 an average hourly wage of \$21, a reasonably large amount on online platforms). We first make sure participants understand the bidding mechanism and its consequences through an attention-demanding quiz. For those that pass the quiz, we then find that, in our main experimental setting, participants integrate social externalities into their pricing bids, even though buying the stock does not change whether the charity transfers happen (see Section 3 for a 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: As we explain in the paper, this is consistent with participants valuing 61c each dollar given to, or taken from, charities by the firm. We find that the scaling of non-pecuniary preferences is close to linear: doubling the size of a social externality doubles its impact on willingness to pay (see Fig. 1, more details below).

To further elicit how much participants care about impact, we construct a second version of our experiment where bids are pivotal, meaning that charity transfers only happen if participants buy the stock. We find that this does not materially change bidding behavior. Therefore, we find *no* evidence of impact-seeking motives. Note here that our participants clearly *understand* the consequences of their bidding behavior: Our practice quiz tests their understanding that their bid is, in this case, pivotal. Participants who do not get a *perfect* score on this quiz are excluded from the experiment.

We then test whether participants conflate prosocial bidding with direct charity donation. A possibility is that participants are confused: They may think that their excess bid goes to the charity. While our quiz is designed to ensure that they have a perfect understanding of all cash movements, participants may have forgotten when they are at the bidding stage. In this case, we would expect that prosocial bidding is a substitute for direct donation. We find two results showing no substitution between donation and prosocial bidding, i.e., that the two behaviors are independent. First, in most conditions, we allow participants to directly donate to the charity. We find donations to be uncorrelated with pro-social bidding in the cross-section of participants. Second, in one of our conditions, participants are *not* allowed to donate, so that we can measure the effect of this option on prosocial bidding. We find that the option to donate does not change the sensitivity of the bid to charity value. Overall, these results are inconsistent with consequentialist altruistic preferences: In the non-pivotal condition, participants are willing to overpay (the experimenter) for charity-giving stocks, but much less willing to directly donate (to charities). Yet, only the second option really has impact.

We finally run a few robustness checks. First, participants recruited through MTurk may not be representative of the population of stock investors. To partially address this, we also run our baseline experiment with participants from a different platform, *Prolific*. This platform allows us to target participants who (1) have at least a college degree

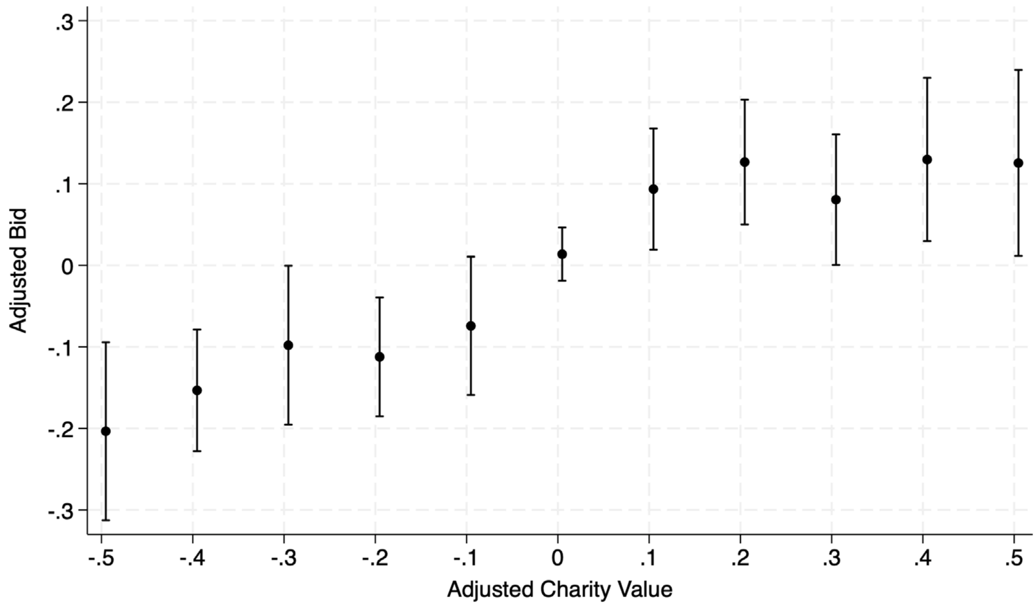

Fig. 1. Linearity of the Relationship between Adjusted Bid and Adjusted Charity Value.
Note: This figure plots the binned scatter plot and 95% confidence intervals of bids against charity values after controlling for selfish values using the "binsreg" stata package developed by Cattaneo et al. (2023). We pool together the "Baseline" and "Baseline with Donation" treatments. We demean the variables before constructing the bin scatter plot.

and (2) claim to be using an online stock trading platform.[1] Another advantage of *Prolific* is that it tends to provide higher quality responses than MTurk (especially since COVID, Douglas et al., 2023), which makes this a natural robustness check to run. When comparing participants from MTurk (in 2019) and Prolific (in 2022 and 2023), we find that treatment effects are statistically indistinguishable, although Prolific survey takers show a tiny glimmer of impact-seekingness. We further explore heterogeneous effects along education, gender, and attentiveness. While gender has no real effect on prosocial bidding and prevalence of deontological behavior, college graduates seem to display some small propensity to seek impact compared to high-school graduates (all of them coming, by design, from MTurk). The effect is only marginally significant ($p = .082$), but a larger experiment (we already have 2496 bidding rounds) may help uncovering it.

Another robustness check that we run tests bidding sensitivity to the size of stakes. In a separate experimental condition (with both pivotal and non-pivotal bidding), we simply multiply stakes by 5. Comparing this with our main sample, we find no difference in prosocial bidding, and no difference in the insensitivity of bidding to pivotality.

The paper is structured as follows. Section 2 connects our findings to three strands of economic literature. Section 3 develops a simple analytical framework that can be used to analyze results. Section 4 describes our experimental design. Section 5 analyzes in detail our main experimental results and their economic interpretation. It also analyzes the substitution with direct donations. Section 6 studies the robustness to various features of the experimental setting. Section 7 concludes.

2. Literature review

Our paper is related to three 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. For instance, Ottoni-Wilhelm et al. (2017) test the presence of deontological preferences by varying stake size. Their idea is that, when a cause already receives enough money from others, consequentialist thinkers would be less likely to donate themselves, since the goal they care about is already reached. In our set-up, we measure consequentialist preferences by directly manipulating impact. 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 a 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 attitudes to reciprocity/fairness) predict the willingness to invest in responsible financial products and to accept higher fees. In particular, by pairing

[1] A limitation of this approach is that most sociodemographics are self-reported. But they are reported *to the survey platform* upon enrollment with Prolific, thus independently of our particular experiment. While participants may still choose misrepresent who they are, their incentives are not distorted by our specific experiment. This somewhat alleviates concerns of experimenter's demand. Another reason to believe that participants truthfully report their characteristics is that, for tasks that can be verified (like attention checks), data quality on this platform tends to be nearly perfect (Douglas et al., 2023).

individual investment data with the results of an experiment, they establish that the decision to invest in a socially responsible fund is primarily linked to pro-social preferences. They measure prosocial preferences via the second-mover response in a trust game. Our results are consistent with their findings that (1) individuals are ready to forego monetary payoffs when investing responsibly, and (2) that responsible investing is not a simple substitute for charity donation. Unlike Riedl and Smeets (2017), we do not have the ability to match real mutual fund investment decisions to our experiment participants. We instead resort to investment in firms that are synthetic but involve real financial payoffs. Our main contribution is to precisely characterize the nature of social preferences by drawing a distinction between impact-driven motivations and value-alignment motivations of our subjects. 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.

Our findings vindicate modelling assumptions in Heinkel et al. (2001) and more recently Pástor et al. (2021) and Pedersen et al. (2021), who develop equilibrium models where a fraction of investors have a distaste for holding firms that are not clean. These models find that "dirty" companies trade at a discount compared to their "clean" peers, because in equilibrium. In line with these findings, 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. (2011) 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 recategorized as having a poor sustainability footprint (even though there is no real change in content). Our contribution to this literature is to isolate the effect of stockholder moral preferences on 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), like us, use 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. However, 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 donations.

In a contemporary paper, Heeb et al. (2022) also evaluate the willingness to pay for shares of funds with varying CO2 footprints. They find that participants are inclined to pay more for investing in sustainable funds. A key difference between their design and ours, is that our stark design allows to clearly separate value alignment from impact. In our setting, there is no ambiguity about whether buying shares has a social consequence or not, and all payoffs are measured in dollars. An advantage of the setting in Heeb et al. (2022) is that it is closer to a real situation where investors consider the CO2 footprint of a portfolio, but the downside is that participants might be unsure of whether and how their investment decision does or does not affect companies' emissions. As a result, it is difficult to know if participants trade for impact or value alignment.

Finally, our results resonate with a political science literature that investigates the moral motivations behind voting. For instance, Federsen et al. (2009) run an experiment where participants are asked to vote about a moral outcome (vs. a selfish one), and vary the extent to which voters are pivotal. They find that participants mostly express deontological preferences. In their paper, as pivotality decreases, voters do not seem to vote more for the ethical alternative (as they should if some of them were consequentialist). In that sense, their experiment lines up with our evidence, which shows no evidence of consequentialist behavior.

3. Model

We model here our experimental condition which uses the BDM bidding mechanism (Becker et al., 1964): Participants bid for a stock that gives part of its profits to charities. In the model, we allow for limited attention. We explore value-alignment and impact-driven preferences, and charity donation. This framework allows us to derive and interpret our simple empirical specification.

3.1. A simple model of investors' social preferences

The model is inspired by Leszczyc and Rothkopf (2010). We start with the non-pivotal condition, where the participant bids for a stock which donates c to the charity, regardless of whether the bid is successful or not. The participant's utility from holding the stock is given by $u_h = s + \alpha c + v_h$, where s is the dividend and c the money paid by the firm to the charity. αc reflects the utility from the charity getting c *when the participant holds the stock*. v_h is an idiosyncratic noise that varies across participants. If the participant does not hold the stock, her utility is $u_r = \beta c + v_r$. βc reflects utility from donation c going to the charity *when she does not hold the stock*. A purely consequentialist participant should only care about the good that is done, independently of whether she holds the stock or not, i.e. $\alpha = \beta$. Alternatively, classic asset pricing models of ESG investing (e.g. Pástor et al. (2021)) assume that a socially responsible investor enjoys a private benefit from holding responsible stocks, i.e. $\alpha > 0$ and $\beta = 0$.

In our experiment, participants participate in a second price auction against a random draw (the BDM bidding mechanism). The participant places her bid b, while another price p is drawn from a uniform distribution with support $[0, \bar{p}]$. The participant only wins the auction if $b > p$, and then pays p. Thus, the expected utility from bidding b is given by:

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

and the optimal bid b^* is given by $b^* = s + (\alpha - \beta)c + v_h - v_r$

Now, we assume that the participant is inattentive. Relying on Gabaix (2019), we assume that the participant chooses the weighted average of a default bid $\bar{b}$ and the optimal bid b^*. Let λ be the weight on the optimal bid, so that the bidding decision writes:

$$b = (1-\lambda)\bar{b} + \lambda s + \lambda(\alpha - \beta)c + \epsilon \quad (1)$$

where we note $\epsilon = \lambda(v_h - v_r)$. $\lambda = 1$ corresponds to the case of full attention.

Since we randomize over s and c, regressing b on s and c identifies λ and $\alpha - \beta$, but not α and β separately. Identifying β is, however, interesting, since it measures the part of investor motivation that is consequentialist. $\beta = 0$ corresponds to the case of pure value-alignment.

To pin down β, we use the pivotal condition, where the charity does not receive any money when the bid fails. In this case, $u_r = v_r$. The observed bid is thus given by:

$$b = (1-\lambda)\bar{b} + \lambda s + \lambda \alpha c + \epsilon \quad (2)$$

Empirical Specification: In our empirical analysis, we run the following regression, for bid i:

$$b_i = \gamma + \delta s_i + \mu c_i + \epsilon_i \quad (3)$$

Our model provides a structural interpretation for this equation. We find that the sensitivity of bids to cash flows s_i to be the same in the pivotal and the non-pivotal condition. This establishes that the attention parameter λ is the same in both conditions. Then, if $\widehat{\mu_P} > \widehat{\mu_{NP}}$, then Eqs. (1) and (2) show that $\beta > 0$, so that there is some consequentialist motive in bidding decisions. If, however, $\mu_P = \mu_{NP}$, then $\beta = 0$. Bidders only care about the donation when they own the stock: They are 100% value aligned. We will find evidence of this (Section 5).

3.2. Donation option

A natural hypothesis is that investors might confuse prosocial bidding for direct donation. In the baseline condition of our set-up (the non-pivotal case), note that this would be an error, since pro-social bidding amounts to transferring money to the seller (the experimenter), not the charity. To clarify this, assume participants can directly donate d at the end of the experiment. To fix ideas, focus on the non-pivotal condition and start from the benchmark model with full attention. Then, utilities are given by:

$$u_h = s + \alpha c + h(d) - d + \epsilon$$

$$u_r = \beta c + h(d) - d + v$$

where h is the utility from the charity receiving d. In this benchmark, we assume that (1) the utility function is separable in charity value and direct donation, and (2) utility from direct donation does not depend on bid success.

This model suggests two tests. First, one can regress observed donation d on the charity value c. If utility is separable, agents consider prosocial bidding and donating as separate problems. Then, it is straightforward to see that optimal donation is uncorrelated with charity value.[2] We will find evidence consistent with such separability.

The second test consists of comparing the bid-to-charity value sensitivity when donation are impossible (the baseline) and when donations are possible. In the above separable model, it is easy to see that excess bids are still given by Eq. (1), whether d is chosen optimally or constrained to zero. We will find evidence consistent with this. Thus a model where utilities u_h and u_r are separable, and the donation part of these utilities does not depend on bid success, will fit the data well.

4. Experimental design

We first describe the overall structure of our experiment. We then discuss how the experiment was implemented across three batches. We recruit participants through Amazon's Mechanical Turk (MTurk) in our first batch, and the platform Prolific in our second and third batches.

4.1. Overall structure

In our experiment, participants bid for shares that vary by how much cash dividend they pay, and by how much money they add to, or remove from, a fund that will be donated to charities. Each participant bids on three different company types in random order. The "ethical" company gives away some shareholder profits to a charity "wallet"; the "unethical" company takes money away from the charity "wallet" and gives it to shareholders; and the "neutral" company neither gives to, nor takes money from, the charity wallet. To participate in the experiment, participants have to pass a practice quiz that ensures they can understand the cash movements between the firm, their own wallet and the charity wallet.

To elicit the truthful valuation of the dividend and the charity donation, we use the classic Becker–DeGroot–Marschak (BDM) bidding mechanism (Becker et al., 1964). 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.3, to which we return below). As discussed in Section 3, the BDM mechanism will allow us to measure the extent of attention and the structure of pro-social bidding.

We define three variables:

1. "Selfish value" is the cash dividend directly paid to the participant by the firm (s in Eq. (3)). 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 (c in Eq. (3)). 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. "Bid" is the amount that is effectively bid by the participant (b in Eq. (3)).

The model in Section 3 provides a natural interpretation for the cross-sectional regression of b on s and c. The coefficient on s allows for measuring imperfect attention. This coefficient should be 1 if attention was perfect. In a recent paper, in the absence of charity donation, Charles et al. (2024) find it to be around 0.6. We will obtain similar orders of magnitude. The coefficient on c measures the combination of imperfect attention and pro-social bidding, with a slightly different interpretation depending on whether bidding does (pivotal condition) or does not affect donation (non-pivotal condition), as explained in Section 3.

4.2. Sequence of events: Baseline condition

We now describe 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, 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 is pledged 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* (or *non-pivotal*) condition, we make it clear that both successful and failed bids lead to the *same changes* in the charity wallet. As discussed in Section 3, purely impact-driven investors should not be

[2] It simply comes from maximizing u_h and u_r with respect to d for given c, since donations are made after bidding.

willing to pay more than s, the selfish value. Even if they are altruistic, consequentialist investors are not expected to bid at a premium for ethical firms (or at a discount for unethical ones), since their action has no consequence for the charity.

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 they fully understand how the bidding game works and the exact consequences of their choices on the various cash movements. 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. Those who do not receive a perfect score fill out demographic questions about their background, and then exit the experiment. They receive the base payment to compensate them for their time.

The details of the practice quiz are 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, we set the parameters of the quiz so that 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 price is set at $0.50, while we tell the participants to consider a hypothetical bid of $1.1, so that the company share is actually purchased at this bid (for a price of $0.50). In the "fail" scenario, the random price is set at $2, above the hypothetical bid value of $1.1, so that the company share is not purchased. Participants see four scenarios drawn at random: The first firm is ethical with probability 1/2 (in which case the second firm is unethical), and the bid is successful with independent 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 at this hypothetical bid and random price, 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 see the three company types (unethical, ethical, neutral) in random order. Each round, the company randomly draws a profit $\pi \in \{0.5, 0.6, 0.7, 0.8, 0.9, 1\}$. If neutral, it gives the entirety of this profit to the shareholder ($s = \pi$). If ethical, the company gives a random portion of this profit to the charity $c \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$, and the rest, $\pi - c = s$, to the shareholder. The unethical company takes money from the charity wallet $-c \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. The participants obtains the profits plus the charity money $\pi - c$. Note that this sequential randomization induces a mechanical correlation between charity donation c and selfish value $s = \pi - c$. This is not a problem for our regression specification, which uses both of these variables as regressors. Finally, the random price generated for the BDM mechanism is between 0 and $2.

We provide a screenshot of the bid for an unethical company in Appendix Figure A.4(a). 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 in either the practice quiz or the main experiment — and find that they are not (See Appendix Table C.2).

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, etc.).

In the next subsection we describe each of the treatment conditions in our three experiment batches.

4.3. Main treatment conditions (batch 1)

We conducted the first batch of experiments on MTurk to test the model in Section 3. There are two separate installments, first on July 22, 2019, and then two weeks later on August 5th, 2019. We have 448 participants in total, of which 369 passed the quiz, allocated to 3 main conditions already discussed in Section 3[3]:

1. *Baseline/non-pivotal:* The pure baseline condition is exactly the one described in the previous section. Participants are asked to bid on three firms, and the charity donation c occurs regardless of whether the individual ends up owning the stock or not. Appendix Figure A.4(a) shows a screenshot of baseline condition for an example unethical firm.
2. *Baseline/non-pivotal with donation*: This condition is the same as the baseline described above, except that participants can directly donate to the charity from their personal wallet. This opportunity is offered at the end of the experiment, but advertised in the first screen. We provide a screenshot of the donation screen in Appendix Figure A.5.
3. *Pivotal with donation*: In the pivotal condition, *no money* will be added or removed from the charity wallet if the participant does not end up buying the stock (see Appendix Figure A.4(b) for a sample screenshot). The practice quiz is similarly adapted, ensuring that participants are fully aware that, if the bid fails, there is no transfer to charity (see Appendix Figure A.6 for a screenshot of an example unethical firm). As in the second condition, donation is an option at the end, and advertised at the beginning (Appendix Figure A.5)

These three conditions directly test the hypotheses outlined in Section 3. Comparing conditions (2) and (3) allows us to estimate α (hedonic flow of owning a stock that donates) and β (hedonic flow of not owning it). Furthermore, comparing conditions (1) and (2) allows us to determine to what extent bidding and direct donations are substitutes.

4.4. Prolific (batch 2)

One possible concern with our MTurk experiment is that participants are not representative of investors. In order to partially address this, we conducted one last round of experiments on the platform Prolific, which has two main advantages over MTurk. First, participants have to register with the platform, and the platform verifies some of the self-reported information. This helps improve the reliability of reported socio-demographics relative to MTurk. Second, the platform allows for selecting participants according to some of these characteristics. We use this second feature to restrict the sample in this last batch to participants who (1) use a stock trading platform and (2) have a degree from 2 or 4 year college. In doing so we are hoping to narrow the focus on college graduates with some financial knowledge, participants who are a bit closer in demographics to the average investor.

[3] There are two additional treatments from the second wave: Stocks purchased on behalf of another participant, and Randomized charity (see Appendix D) Although including these two treatments would not change our main results, we omit them from the paper because they do not directly test our core hypothesis, and would therefore reduce the clarity of exposition.

We ran this second batch in only one round on July 6, 2022, and recruited 300 participants, of which 250 passed the quiz. It contains two conditions, to which participants are allocated randomly: (1) Baseline with Donation and (2) Pivotal with Donation. These conditions are *identical in all respects* to the two conditions with donation in first batch. In both cases, and as in the first batch, participants have the option to donate part of their earnings at the end of the experiment, and are told so in advance.

4.5. Prolific & symmetry (batch 3)

This last batch reproduces the experiments of the second batch, with more symmetric wording in the non-pivotal condition. A possible worry in batches 1 and 2 is that, in the non-pivotal condition, the wording is slightly asymmetric between successful and unsuccessful bids: We do not say explicitly in the prompt what happens when the bid fails. There is no such asymmetry in the pivotal condition (see Appendix Figure A.4) To check that our results were not driven by this, we ran, on May 2023, a slightly modified version of the experiment in Batch 2, where the only difference is that we add "if you do not buy the share" when the bid fails in the non-pivotal condition (see Appendix Figure A.7) This is a bit more explicit.

This last batch of experiments has 299 participants, of which 213 passed the quiz. The conditions are otherwise identical to Batch 2. We first checked whether our main specification yields different results in this new batch. We show in Appendix Table C.1 that they are not. The sensitivities of bids to charity value and selfish values are not significantly different across conditions. They are strongly significant in Batch 2, where the slightly smaller t-stats can be explained by the fact that Batch 3 is smaller than Batches 1 and 2 taken together. Overall, given that our regression yields the same result in Batch 3 as in the previous two batches, in spite of the fact that the wording slightly differs, we decide to pool all three batches to maximize statistical power.

4.6. Robustness to larger stakes

At the same time as our third batch of participants, we also ran a separate condition to test the effect of larger stakes (in May 2023 as well). This condition is exactly identical to Batch 3, except that stakes are multiplied by 5. Profits are drawn from $\{2.5, 3, 3.5, 4, 4.5\}$. Charity values are drawn in $\{0.5, 1, 1.5, 2, 2.5\}$. Because the stakes are five times larger in these conditions, we do *not* pool it with our main tests, but only use this sample for robustness. We later directly compare regression results in our pooled sample (batches 1,2,3) and results in this "larger stakes" condition (this will be in Table 9, described below). As we will see, pooling large stakes with our 3 batches will bias our econometric specification because large stakes have a larger constant (and would thus bias pro-social bidding upwards).

In total, 298 participants to the third batch are randomly assigned to this "larger stakes" condition. Out of these, 245 passed the quiz: 116 in the non-pivotal condition, 129 in the pivotal condition. These conditions are exactly the same as Batch 3: the quiz that participants must pass depends on pivotality, participants are recruited on Prolific, the set-up is identical to Batch 3 (with symmetric wording), and there is a possibility to donate at the end (see the fourth column of 1).

4.7. Descriptive statistics

We describe here our main sample of Batches 1,2 and 3 (and do not discuss the "Larger Stakes" sample, which we will analyze separately). Table 1 shows summary statistics about quiz passing rates, payments and condition characteristics across the three batches and five waves of our experiment. There are 1047 participants in total, but only 832 passed the quiz (passing rate: 79%). The number of participants per condition is around 110, slightly larger than this in our first batch, slightly smaller in the last batch. Generally speaking, participants in Batch 3 tended to fail the quiz more often, and be slower, but not significantly so. Bonuses hover around \$2.2 and charity donations around \$1.1. Overall, our pooled sample has 2496 bids because each of the 832 participants bids on the three company types: (1/3 neutral, 1/3 ethical, and 1/3 unethical firms).

Overall, demographic groups are reasonably balanced across MTurk treatments (allocation is random, but we ran several waves over a period of 4 years). Table 2 presents the summary statistics on demographics by treatment condition and platform. Panel A presents statistics on participants in the main experiment on MTurk, and Panel B presents statistics on participants in the main experiment on Prolific. About 60% of the participants are male. Consistent with the notion that Prolific allows to filter for college degree and stock trading, albeit imperfectly, we find that, among those who pass the quiz, the shares of college graduates and stock owners both increase dramatically. Clearly, the Prolific filters do not work perfectly (share of those who report trading stocks is only 60% in our survey question), but they seem to shift the composition of the pool of participants in a significant way.

Table 1
Description of the experiment and treatments.

	Main sample			Robustness
	Batch 1	Batch 2	Batch 3	Larger stakes
Panel A: Description				
Platform	Mturk	Prolific	Prolific	Prolific
Experiment Year	2019	2022	2023	2023
Participants (count)	448	300	299	298
Passed Quiz (count)	369	250	213	245
Quiz passing rate (%)	82.37	83.33	71.24	82.21
Average duration (mins)	11.76	14.76	21.03	15.11
Base Payment (\$)	2.00	3.00	3.00	3.00
Participant's wallet (average \$)	2.31	2.05	2.19	10.01
Charity wallet (average \$)	1.05	1.23	1.28	5.96
Panel B: Main experiment treatment conditions (Number of participants)				
Baseline	114	–	–	–
Baseline with donation	131	118	124	116
Pivotal with donation	124	132	89	129

Table 2
Demographics by treatment condition and platform.

	Baseline	Baseline with Donation	Pivotal with donation
Panel A: MTurk			
18–35	0.54	0.56	0.56
35–55	0.32	0.35	0.40
55+	0.13	0.08	0.04
College	0.68	0.63	0.51
Male	0.61	0.57	0.62
Fin. Lit.	0.76	0.78	0.81
Stocks	0.29	0.25	0.19
Panel B: Prolific			
18–35	0.40	0.44	
35–55	0.44	0.40	
55+	0.16	0.16	
College	0.97	1.00	
Male	0.57	0.67	
Fin. Lit.	0.88	0.84	
Stocks	0.59	0.64	

5. Main results

In this section, we first explore the baseline/non-pivotal condition. We then compare this baseline with the pivotal condition, which allows for separately identifying the parameters α and β discussed in Section 3, i.e. the share of deontological and consequentialist motivations in participants' bids.

Table 3
The sensitivity of bids to charity values: Baseline.

	LHS: Bid		
	(1)	(2)	(3)
Charity value	.37***	.37***	
	(.073)	(.073)	
$\frac{1}{2}$ Charity value^2		−.4*	
		(.24)	
(Charity value)$^-$			.46***
			(.091)
(Charity value)$^+$			.28***
			(.086)
Selfish value	.61***	.62***	.62***
	(.069)	(.069)	(.069)
Constant	.28***	.29***	.29***
	(.055)	(.055)	(.055)
Observations	1,461	1,461	1,461
R^2	0.08	0.08	0.08
P-value asymmetry			.082

Note: The LHS is the bid, b. We regress the bid on the dividend s (Selfish value), and the amount transferred by the charity to the company c (Charity value). We pool together all baseline-related conditions (baseline and baseline with donation). We report in the last line the p-value is for the t-test of equality of coefficients for $(Charityvalue)^-$ and $(Charityvalue)^+$. Standard errors in parentheses, clustered at the worker level. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table 4
The sensitivity of bids to pivotality.

	LHS: Bid		
	(1)	(2)	(3)
	Pooled	Baseline	Pivotal
Charity value	.45***	.37***	.55***
	(.057)	(.073)	(.089)
Selfish value	.63***	.61***	.65***
	(.053)	(.069)	(.081)
Constant	.25***	.28***	.22***
	(.042)	(.055)	(.064)
Observations	2,496	1,461	1,035
R^2	0.07	0.08	0.07
P-value of equality (Charity value)		.127	
P-value of equality (Selfish value)		.738	

Note: The LHS is the bid, b. We regress the bid on the dividend s (Selfish value), and the amount transferred by the charity to the company c (Charity value). Column 2 uses the "baseline" conditions from Table 3; column 3 uses the "pivotal with donation" condition, and Column 1 include the pooled estimates. We report in the second-to-last line the p-value of the t-test that the regression coefficients on "Charity value" for the "baseline" treatment is equal to the coefficient for the "pivotal with donation" treatment. We report in the last line the p-value of the t-test that the regression coefficient on "Selfish value" for the "baseline" treatment is equal to the coefficient for the "pivotal with donation" treatment. Standard errors in parentheses, clustered at the worker level. * p < 0.10, ** p < 0.05, *** p < 0.01.

5.1. Baseline: Evidence of value alignment

We start by restricting our focus to participants in the baseline conditions, where bids are non-pivotal. Across all three batches, this leads to 487 participants, or 1,461 rounds of bidding. Note that this approach puts together 373 participants who were allowed to directly donate, and 114 who were not. As we will see below, estimates do not significantly differ between the types, so we pool them all to gain power.

As discussed in Section 3, our empirical strategy consists of running the following regression:

$$b_i = \gamma + \delta s_i + \mu c_i + \varepsilon_i$$

where δ measures the degree of attention to the bidding process, while μ conflates the effect of attention and non-consequentialist altruistic considerations in bidding. μ is expected to be zero if, in the model, $\alpha = \beta$: if participants are impact-driven enjoy the warm glow of charity donation, whether they own the stock or not. If however owning the stock makes a difference, $\mu > 0$ ($\alpha > \beta$), and participants have some deontological preference.

In Fig. 1, we provide graphical evidence on the sensitivity of bids to charity donation. We show a binned scatter plot, controlling for selfish value s. We need to control for s because, as explained above, our randomization design induces a mechanical correlation between c and s. To do this, we use the approach recently developed by Cattaneo et al. (2023). As it appears quite clearly, there is a positive, monotonic relationship between bids and charity value, after controlling for s. Overall, this is evidence that value alignment matters: Participants bid higher for companies that donate more (even though their bid has no impact).

Table 3 shows the formal regression. Again, here, we focus on the same 1461 rounds of bidding of the 3 quasi-baselines. Consistent with Fig. 1, we see in column 1 that the sensitivity of bidding to charity value, controlling for selfish value, is .37 ($t = 5.1$), i.e., for each $1 increase in charity donation, participants increase their bids by 37c. As argued in Section 3, this coefficient conflates the effect of attention to bidding and participant altruistic concern. In order to disentangle the two, we can use the coefficient on s, which is purely driven by attention (under full attention, the coefficient on s should be 1). Looking in Column 1, we see that the sensitivity of bidding to s is .61 ($t = 8.8$), consistent with the idea that attention is significant, but not complete (Charles et al. (2024) find a similar level of attention in their related experimental set-up). We show in Appendix Figure B.1 that the relationship between bidding and s is close to linear. As a result, the estimate of the "net deontological preference" $\alpha - \beta$ is given by $.37/.61 = .61$: For each $1 of charity donation, the stock-owning utility of the participant increases by 61c. This is a large effect, presumably due to the starkness of our experimental set-up.

Columns 2 and 3 test for non-linearity. Consistent with Fig. 1, both columns show some limited evidence of concavity. In column 2, a squared term in s is significant at 10%. In column 3, negative charity donations reduce bidding by 46c per dollar, while positive charity donations increase it by only 28c. The difference is however, only significant at 8%. There is a small indication that prosocial bidding is asymmetric, consistent with some form of social loss aversion. But since this effect is weak, we focus on the linear specification in the remainder of the paper.

5.2. Pivotal condition: Evidence of no impact-seeking behavior

Our baseline condition only identifies $\alpha - \beta \approx .61$, as discussed in Section 3. It indicates that value alignment is an important source of utility (at least .61), but does not prove that impact-seeking behavior is absent: It could still be that $\beta > 0$, i.e. that participants care about the charity donation even when they do not own the stock.

To separately identify α and β, one needs to compare the pivotal and the non-pivotal conditions. To the 1461 non-pivotal bids studied so far, we thus add the 1035 pivotal bids collected in batches 1, 2 and 3. We show in Fig. 2 a binned scatter plot of bids against charity values for the two conditions separately. As in Fig. 1, these plots control for the selfish value s, which is a separate driver of the cross-section of bids and negatively correlated (by design) with charity donations. As it clearly appears in the figure, the slopes of pivotal and non-pivotal conditions are rather similar.

Table 4 runs formal statistical tests. In all columns, we estimate our empirical specification. Column 1 uses the entire, pooled sample (all of batches 1,2 and 3). Column 2 focuses on non-pivotal (baseline) bids, and column 3 on pivotal bids. First focusing on the sensitivity to selfish value, we find .61 in the baseline vs. .65 in the pivotal conditions. The difference in attention is insignificant (p value = .74). Given that

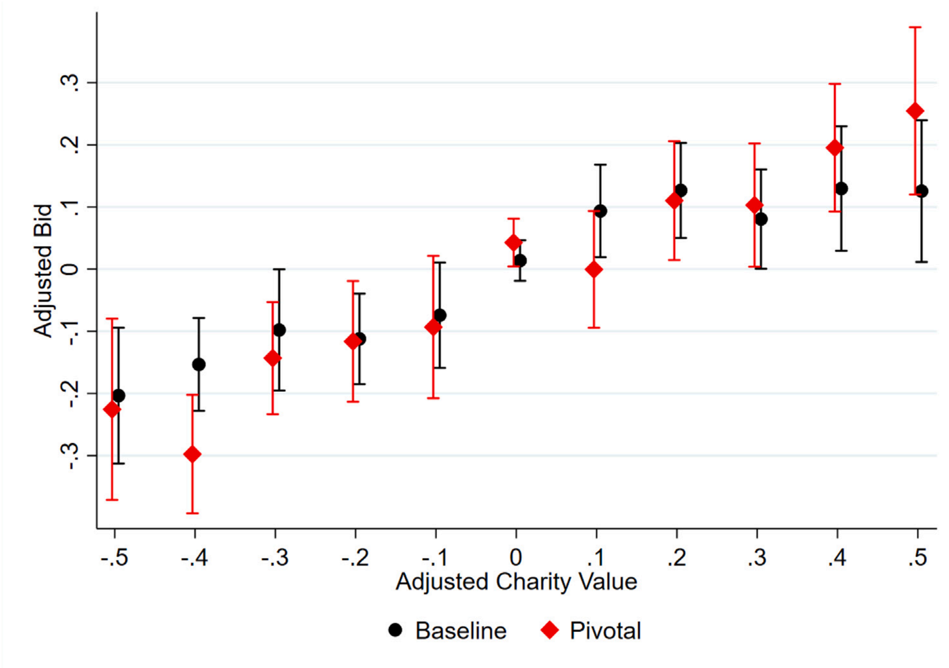

Fig. 2. Adjusted Bids and Adjusted Charity Values: Pivotal v Baseline.
Note: This figure plots the binned scatter plot and 95% confidence intervals of bids against charity values after controlling for selfish values using the "binsreg" stata package developed by Cattaneo et al. (2023) The black chart is executed on the "Baseline" treatments (same as in Fig. 1) and the red chart on the "Pivotal" condition. We demean the variables before constructing the bin scatter plot.

Table 5
Effect of corporate externality on propensity to donate.

	(1)	(2)
	Y/N	Amount
Total charity value	−.07	.32
	(.084)	(.28)
Constant	.24***	.82***
	(.016)	(.058)
Observations	718	174
R^2	0.00	0.01

Note: The data used here corresponds to the two conditions with the option to donate: Baseline with donation and pivotal with donation. 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. Standard errors in parentheses, clustered at the worker level. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table 6
Sensitivity to option to donate.

	LHS: Bid	
	(1)	(2)
	Baseline	Baseline with donation
Charity value	.51***	.33***
	(.14)	(.084)
Selfish value	.86***	.54***
	(.13)	(.08)
Constant	.11	.33***
	(.091)	(.065)
Observations	342	1,119
R^2	0.15	0.06
P-value of equality (Charity value)	.266	
P-value of equality (Selfish value)	.029	

Note: The LHS is the bid, b. We regress the bid on the dividend s (Selfish value), and the amount transferred by the charity to the company c (Charity value). Column 1 uses the *baseline* condition (with charity choice and no donation) and column 2 uses the *baseline with donation* condition (with charity choice and the donation option). We report in the second-to-last line the p-value of the t-test that the regression coefficient on "Charity value" for the "baseline" treatment is equal to the coefficient for the "baseline with donation" treatment. We report in the last line the p-value of the t-test that the regression coefficient on "Selfish value" for the "baseline" treatment is equal to the coefficient for the "baseline with donation" treatment. Standard errors in parentheses, clustered at the worker level. * p < 0.10, ** p < 0.05, *** p < 0.01.

attention is the same in both conditions, we can directly attribute the difference in sensitivities to c to consequentialist preferences ($\alpha \neq 0$). As previously shown in Table 3, the sensitivity in the non-pivotal condition is .37 ($t = 5.1$), versus .55 ($t = 6.2$) in the pivotal condition. Both are strongly significant, but statistically indistinguishable with $p = .13$. Overall, while there is evidence that $\alpha > 0$ (i.e. investors care about charity donation even when they do not own the stock), we cannot reject the hypothesis that $\alpha = 0$ at conventional levels of significance. Our participants are close to being 100% deontological.

5.3. Substitution between personal and corporate donation

A natural question is whether our participants confuse prosocial bidding with generosity. For instance, by paying a higher price, they may be under the impression that they are actually donating the excess amount (above fundamental value) to the charity — while in reality they are paying the experimenter. We argue that confusion is unlikely, given that participants in the main experiment have to ace the practice quiz. That said, in this section, we show directly that they are not confused and that prosocial bidding is indeed an expression of their deontological preferences.

In Table 5, we focus on the 718 participants who can make donations, and test if donations respond to the total charity value, which varies randomly. The RHS variable, *Total charity value*, is the sum of all transfers to charity in all three rounds experienced by the participants (thus, the sum of the donations of the ethical and unethical firms, since the neutral firm does not donate). Under the null hypothesis of the

Table 7
Robustness to Prolific.

	MTurk			Prolific		
	(1)	(2)	(3)	(4)	(5)	(6)
	Pooled	Baseline	Pivotal	Pooled	Baseline	Pivotal
Charity value	.57***	.51***	.65***	.37***	.27***	.49***
	(.087)	(.1)	(.16)	(.074)	(.1)	(.11)
Selfish value	.77***	.78***	.76***	.54***	.5***	.59***
	(.076)	(.093)	(.13)	(.07)	(.097)	(.1)
Constant	.14**	.13*	.14	.33***	.39***	.26***
	(.06)	(.073)	(.11)	(.056)	(.078)	(.079)
Observations	1,107	735	372	1,389	726	663
R^2	0.10	0.11	0.09	0.06	0.06	0.06
P-value of equality (Charity value)		.464			.138	
P-value of equality (Selfish value)		.9			.507	

Note: In all regressions, the LHS is the bid, b. We regress the bid on the dividend s (Selfish value), and the amount transferred by the charity to the company c (Charity value). For Columns (1)–(3), we use the data from the 2019 Amazon MTurk implementation, and Columns (4)–(6) show the data from the 2022 and 2023 Prolific implementations. Columns (2) and (5) shows the baseline (nonpivotal) treatment. Columns (3) and (6) shows the pivotal treatment. and Columns (1) and (3) shows the nonpivotal and pivotal treatments combined. We report in the second-to-last line the p-value of the t-test that the regression coefficient on "Charity value" for the "baseline" treatment is equal to the coefficient for the "pivotal" treatment. We conduct this separately for the MTurk sample and for the Prolific sample. We do the same t-test in the last line for the coefficient on Selfish value. Standard errors in parentheses, clustered at the worker level. * p < 0.10, ** p < 0.05, *** p < 0.01.

model in Section 3, corporate donation and participant donation are separable and additive, so that the expected coefficient is zero.

In Column (1), the LHS variable is a dummy equal to 1 if the participant donates at all, and zero otherwise. In column (2), we look at the intensive margin: The amount donated conditional on donating (hence the smaller number of observations). In both cases, the correlation with total firm prosocial behavior (Total Charity Value) is insignificant. Thus, the WTP for virtuous stocks is *not* crowded out by the possibility to donate directly.

Table 6 is a second test of this potential substitution, which is based on comparing the amount of prosocial bidding in conditions where donation is allowed, or not. Columns 1 and 2 thus compare the estimates of Eq. (3) in conditions where participants are allowed to donate, and conditions where they are not. To make things comparable, we focus on the non-pivotal/baseline condition, since this is the only bidding condition for which we have a configuration without donation. Column 1 has 114 participants (baseline, no donation), while column 2 has 373 (baseline, donation). Prosocial bidding is strong and statistically significant in both conditions, but the difference in sensitivity to the charity donation is insignificant (p-value=.27). This confirms results from Table 5 that participants view the decision to donate and the decision to "reward" pro-social companies as unrelated to one another.

6. Robustness

6.1. Sample splits

We first test the robustness of our results by comparing results in MTurk (batch 1) to results in Prolific (Batches 2 and 3). A possible concern with our MTurk participants is that they are not representative of the population of investors. They tend to be less wealthy and less well financially educated. To partially address this, we use evidence from Prolific, a platform that allows to restrict survey participation to subcategories of people. We focus on (1) college graduates and (2) people who have access to a stock trading platform. None of these two demographics are actually verified by the company, but we are hoping the fraction of college graduates and stock traders is larger in our Prolific samples than in our MTurk batch. Table 7 reports the results of our main specification on both samples separately, focusing on pooled conditions (columns 1 and 4), baseline conditions (columns 2 and 5) and pivotal conditions (columns 3 and 6). In spite of the fact that reported college education and stock ownership are more prevalent in Prolific (see Table 2), both subgroups yield the same result: prosocial bidding is strongly significant and large for both pivotal and non-pivotal conditions, and the difference is statistically insignificant, indicating almost pure deontological preferences. We note that Prolific participants tend to be slightly more sensitive to impact, however, suggesting that education and stock ownership may tilt (in an insignificant way) towards consequentialism.

We explore further sample splits in Table 8. Panel A focuses on self-reported education. Consistent with our MTurk/Prolific split, we find some limited evidence that education tilts individuals towards more consequentialism. While prosocial bidding does not depend on pivotality for high-school graduates, all of them from Mturk, ($p = .896$), there is a marginally significant difference among college graduates, coming both from MTurk and Prolific (charity sensitivity of .55 vs .33, with $p = .082$). Panel B splits by gender, and we find no difference between male and female participants: Both exhibit pure deontological preferences. Finally, panel C splits by number of practice quiz attempts, as a proxy for attention. Consistent with expectations, the sensitivity to selfish value, which measures attention in our model (it should be 1 if participants are fully attentive), is higher for participants who passed the quiz in the first attempt. The difference is, however, statistically insignificant. Consistent with stronger attention, prosocial bidding appears a bit stronger (but again not significantly so) for the most attentive participants. In both cases, there is no real difference between the pivotal and non-pivotal treatments.

Overall, sample splits confirm the prevalence of value alignment across sub-populations. There is slight evidence of some impact-seeking motive among more educated participants, but even there, the difference between pivotal and non-pivotal conditions is small.

6.2. Larger stakes

In this last robustness check, we explore the impact of larger stakes. We show the results in Table 9, and report our main results on our pooled sample in columns 1–3 (the same ones as in Table 4). In columns 4–6, we rerun the same regressions, but on the separate sample where stakes are multiplied by 5 (see Section 4 for a detailed description). The table makes it clear that there is no difference in prosocial bidding between regular and large stakes. The sensitivity of bidding to charity is hovers between 0.4 and 0.5, and is not statistically different between non-pivotal and pivotal conditions. The sensitivity to selfish value, which captures equation in Eqs. (1) and (2) is also similar (around .6), indicating that the level of attention also does not depend on the size of stakes.

There is, however, a large difference in the constant, and this justifies our choice to not pool larger stakes with batches 1–3 in our

Table 8
Heterogeneity by characteristics.

Panel A: Education	High School or Less			College or Higher		
	(1)	(2)	(3)	(4)	(5)	(6)
	Pooled	Baseline	Pivotal	Pooled	Baseline	Pivotal
Charity value	.59***	.59***	.56***	.42***	.33***	.55***
	(.14)	(.2)	(.17)	(.062)	(.077)	(.1)
Selfish value	.78***	.84***	.68***	.6***	.57***	.65***
	(.13)	(.18)	(.16)	(.057)	(.074)	(.091)
Constant	.11	.071	.18	.28***	.32***	.22***
	(.1)	(.15)	(.12)	(.045)	(.059)	(.071)
Observations	462	276	186	2,034	1,185	849
R^2	0.09	0.10	0.09	0.07	0.08	0.06
P-value of equality (Charity value)		.896			.082	
P-value of equality (Selfish value)		.527			.526	
Panel B: Gender	Male			Female		
	(1)	(2)	(3)	(4)	(5)	(6)
	Pooled	Baseline	Pivotal	Pooled	Baseline	Pivotal
Charity value	.45***	.37***	.55***	.45***	.37***	.58***
	(.073)	(.097)	(.11)	(.09)	(.11)	(.16)
Selfish value	.69***	.67***	.72***	.54***	.54***	.54***
	(.067)	(.092)	(.098)	(.084)	(.1)	(.14)
Constant	.22***	.25***	.19**	.3***	.32***	.26**
	(.053)	(.074)	(.077)	(.067)	(.082)	(.11)
Observations	1,515	843	672	981	618	363
R^2	0.09	0.10	0.09	0.05	0.05	0.04
P-value of equality (Charity value)		.236			.272	
P-value of equality (Selfish value)		.714			.998	
Panel C: Practice Quiz Attempts	Passed in First Attempt			Passed in Subsequent Attempt		
	(1)	(2)	(3)	(4)	(5)	(6)
	Pooled	Baseline	Pivotal	Pooled	Baseline	Pivotal
Charity value	.53***	.47***	.58***	.4***	.32***	.52***
	(.087)	(.12)	(.12)	(.074)	(.091)	(.13)
Selfish value	.69***	.7***	.69***	.59***	.57***	.61***
	(.079)	(.12)	(.11)	(.07)	(.086)	(.12)
Constant	.19***	.2**	.17**	.29***	.32***	.25***
	(.062)	(.09)	(.082)	(.056)	(.069)	(.093)
Observations	906	477	429	1,590	984	606
R^2	0.09	0.09	0.09	0.06	0.07	0.06
P-value of equality (Charity value)		.52			.202	
P-value of equality (Selfish value)		.96			.766	

Note: In all regressions, the LHS is the bid, b. We regress the bid on the dividend s (Selfish value), and the amount transferred by the charity to the company c (Charity value). All three panels use the main experiment data (Batches 1–3). Columns (1) and (4) shows the nonpivotal and pivotal treatments combined. Columns (2) and (5) show the nonpivotal treatment, and Columns (3) and (6) show the pivotal treatments. We split the sample to show how the pass-through varies by education (Panel A), gender (Panel B), and practice quiz attempts (Panel C). We report in the second-to-last line of each panel the p-value of the t-test that the regression coefficient on "Charity value" for the "baseline" treatment is equal to the coefficient for the "pivotal" treatment. We do the same t-test in the last line for the coefficient on Selfish value. Standard errors are in parentheses, clustered at the worker level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

analysis. Going back to Eqs. (1) and (2), the constant of the regression is equal to $(1-\lambda)\bar{b}$, where λ is the attention parameter and $\bar{b}$ is a scaling parameter that is left unmodeled. Given similar λ in columns 1–3 and 4–6, evidence from the Table suggests that $\bar{b}$ is much larger (4–5 times) when stakes are larger. This is expected, as the participants' anchor should be proportional to the order of magnitudes of payoffs, which are five times larger in the large stake condition. While this difference in the constant can be understood, it prevents us from running a single regression on a completely pooled sample made of the three batches and the large stakes condition. It would induce an upward bias in the coefficients, as c and s are positively correlated with the constant in the super pooled sample.

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 an online experiment that allows participants to submit bids for companies that vary by how much money they add to, 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 a large portion of this charity externality in their bids. More importantly, participants exhibit value-alignment, in the sense that they only care about charity donations when they own the stock, and not if they don't. We find little evidence of impact-seeking motives.

To further learn about investors' moral preferences, our experimental set-up could be extended in several directions. For instance, one could explore the determinants of deontological preferences: it could be that some causes, or other forms of corporate externalities (say, pollution instead of donations), are more likely to trigger consequentialist thinking. Also, one could ask how much the (non-consequentialist) willingness to pay depends on the efficiency of charities receiving money. We leave these leads for future research.

Table 9
Robustness to larger stakes.

	Regular stakes			Larger stakes		
	(1)	(2)	(3)	(4)	(5)	(6)
	Pooled	Baseline	Pivotal	Pooled	Baseline	Pivotal
Charity value	.45***	.37***	.55***	.44***	.3*	.56***
	(.057)	(.073)	(.089)	(.099)	(.16)	(.11)
Selfish value	.63***	.61***	.65***	.64***	.53***	.72***
	(.053)	(.069)	(.081)	(.086)	(.14)	(.093)
Constant	.25***	.28***	.22***	1.1***	1.8***	.5
	(.042)	(.055)	(.064)	(.33)	(.54)	(.34)
Observations	2,496	1,461	1,035	735	348	387
R^2	0.07	0.08	0.07	0.09	0.07	0.14
P-value of equality (Charity value)		.126			.175	
P-value of equality (Selfish value)		.749			.253	

Note: In all regressions, the LHS is the bid, b. We regress the bid on the dividend s (Selfish value), and the amount transferred by the charity to the company c (Charity value).For Columns (1)–(3), we use the main experiment data used throughout the paper. Columns (4)–(6) use the data from the May 2023 treatment with larger stakes, where all parameters were multiplied by 5. Columns (2) and (5) shows the baseline (nonpivotal) treatment. Columns (3) and (6) shows the pivotal treatment. and Columns (1) and (3) shows the nonpivotal and pivotal treatments combined. We report in the second-to-last line the p-value of the t-test that the regression coefficient on "Charity value" for the "baseline" treatment is equal to the coefficient for the "pivotal" treatment. We conduct this separately for the main sample and for the 'large stakes' robustness sample. We do the same t-test in the last line for the coefficient on Selfish value. Standard errors in parentheses, clustered at the worker level. * p < 0.10, ** p < 0.05, *** p < 0.01.

CRediT authorship contribution statement

Jean-François Bonnefon: Writing – original draft. **Augustin Landier:** Writing – original draft. **Parinitha Sastry:** Writing – original draft. **David Thesmar:** Writing – original draft.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

Landier acknowledges research funding from Fondation Bank de France and ADEME.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jfineco.2024.103955.

Data availability

Moral Preferences of Investors : Experimental Evidence -- Replication Package (Original data) (Mendeley Data)

References

Barber, Brad, Morse, Adair, Asuda, Ayako, 2020. Impact investing. J. Financ. Econ..
Bartling, Björn, Weber, Roberto A., Yao, Lan, 2014. Do markets erode social responsibility? Q. J. Econ. 130 (1), 219–266.
Becker, Gordon M., Degroot, Morris M., Marshak, Jason, 1964. Measuring utility by a single-response sequential method. Behav. Sci. 9 (3).
Berk, Jonathan, van Binsbergen, Jules H., 2021. The impact of impact investing. Available at SSRN 3909166.
Bolton, Patrick, Li, Tao, Ravina, Enrichetta, Rosenthal, Howard L, 2019. Investor Ideology. Technical Report, National Bureau of Economic Research.
Broccardo, Eleonora, Hart, Oliver D., Zingales, Luigi, 2020. Exit vs. Voice. Technical Report, National Bureau of Economic Research.
Brodback, Daniel, Guenster, Nadja, Pouget, Sébastien, 2019. The valuation of corporate social responsiblity: An experiment. SSRN Working Paper 4260824.
Cattaneo, Matias D., Crump, Richard K., Farrell, Max H., Feng, Yingjie, 2023. On binscatter.
Charles, Constantin, Frydman, Cary, Kilic, Mete, 2024. Insensitive investors. J. Finance 79 (4), 2473–2503.
Crockett, Molly J., Siegel, Jenifer Z., Kurth-Nelson, Zeb, Dayan, Peter, Dolan, Raymond J., 2017. Moral transgressions corrupt neural representations of value. Nat. Neurosci. 20 (6), 879–885.
Derwall, Jeroen, Koedijk, Kees, Ter Horst, Jenke, 2011. A tale of values-driven and profit-seeking social investors. J. Bank. Financ. 35 (8), 2137–2147.
Douglas, Benjamin, Ewell, Patrick, Brauer, Markus, 2023. Data quality in online human subjects research: Comparisons between MTurk, Prolific, CloudResearch, qualtrics, and SONA. PLoS One 1–17.
Edmans, Alex, 2011. Does the stock market fully value intangibles? Employee satisfaction and equity prices. J. Financ. Econ. 101 (3), 621–640.
Elfenbein, Daniel W., McManus, Brian, 2010. A greater price for a greater good? Evidence that consumers pay more for charity-linked products. Am. Econ. J.: Econ. Policy 2 (2), 28–60.
Falk, Armin, Szech, Nora, 2013. Morals and markets. Science 340 (6133), 707–711.
Federsen, Thimothy, Gailmard, Sean, Sadroni, Alvaro, 2009. Moral bias in large elections: Theory and experimental evidence. Am. Political Sci. Rev.
Gabaix, Xavier, 2019. Behavioral inattention. In: Handbook of Behavioral Economics: Applications and Foundations 1, vol. 2, Elsevier, pp. 261–343.
Gollier, Christian, Pouget, Sébastien, 2014. The "washing machine": Investment strategies and corporate behavior with socially responsible investors. TSE Working Paper.
Green, Daniel, Roth, Benjamin, 2021. The allocation of socially responsible capital. Available at SSRN 3737772.
Hart, Oliver, Zingales, Luigi, 2017. Companies should maximize shareholder welfare not market value. J. Law Finance Account. 2 (2), 247–274.
Hartzmark, Samuel M., Sussman, Abigail B., 2019. Do investors value sustainability? A natural experiment examining ranking and fund flows. J. Finance.
Heeb, Florian, Kölbel, Julian F., Paetzold, Falko, Zeisberger, Stefan, 2022. Do investors care about impact? Available at SSRN 3765659.
Heinkel, Robert, Kraus, Alan, Zechner, Josef, 2001. The effect of green investment on corporate behavior. J. Financ. Quant. Anal. 36 (4), 431–449.
Hong, Harrison, Kacperczyk, Marcin, 2009. The price of sin: The effects of social norms on markets. J. Financ. Econ. 93 (1), 15–36.
Humphrey, Jacquelyn, Kogan, Shimon, Sagi, Jacob, Starks, Laura, 2021. The Asymmetry in Responsible Investing Preferences. Technical Report, National Bureau of Economic Research.
Krüger, Philipp, 2015. Corporate goodness and shareholder wealth. J. Financ. Econ. 115 (2), 304–329.
Landier, Augustin, Lovo, Stefano, 2020. ESG investing: How to optimize impact? HEC Paris Research Paper No. FIN-2020-1363.
Leszczyc, Peter T.L. Popkowski, Rothkopf, Michael H., 2010. Charitable motives and bidding in charity auctions. Manage. Sci. 56 (3), 399–413.
Margolis, Joshua D., Elfenbein, Hillary Anger, Walsh, James P., 2011. Does it pay to be good? A meta-analysis and redirection of research on the relationship between corporate social and financial performance. Research Paper available on SSRN.
Oehmke, Martin, Opp, Marcus M., 2020. A theory of socially responsible investment. CEPR Discussion Paper No. DP14351.
Ottoni-Wilhelm, Mark, Vesterlund, Lise, Xie, Huan, 2017. Why do people give? Testing pure and impure altruism. Am. Econ. Rev. 107 (11), 3617–3633.

Pástor, Luboš, Stambaugh, Robert F., Taylor, Lucian A., 2021. Sustainable investing in equilibrium. J. Financ. Econ. 142 (2), 550–571.

Pedersen, Lasse Heje, Fitzgibbons, Shaun, Pomorski, Lukasz, 2021. Responsible investing: The ESG-efficient frontier. J. Financ. Econ. 142 (2), 572–597.

Riedl, Arno, Smeets, Paul, 2017. Why do investors hold socially responsible mutual funds? J. Finance 72 (6), 2505–2550.

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

Sandel, Michael J., 2012. What Money Can't Buy: The Moral Limits of Markets. Macmillan.

Singer, Peter, 2016. The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically. Yale University Press.

Tasimi, Arber, Gross, James J., 2020. The dilemma of dirty money. Curr. irections Psychol. Sci.

Tasimi, Arber, Wynn, Karen, 2016. Costly rejection of wrongdoers by infants and children. Cognition 151, 76–79.

Zerbib, Olivier David, 2020. A sustainable capital asset pricing model (S-CAPM): Evidence from green investing and sin stock exclusion. Available at SSRN 3455090.