



Testing the Motivational Strength of Positive and Negative Duty Arguments Regarding Global Poverty

Luke Buckland¹ · Matthew Lindauer²  · David Rodríguez-Arias³ · Carissa Véliz⁴

Accepted: 19 April 2021 / Published online: 13 May 2021

© The Author(s) 2021

Abstract

Two main types of philosophical arguments have been given in support of the claim that the citizens of affluent societies have stringent moral duties to aid the global poor: “positive duty” arguments based on the notion of beneficence and “negative duty” arguments based on noninterference. Peter Singer’s positive duty argument (Singer 1972) and Thomas Pogge’s negative duty argument (Pogge 2002) are among the most prominent examples. Philosophers have made speculative claims about the relative effectiveness of these arguments in promoting attitudes and behaviors that could lead to the alleviation of poverty. In this article we present the results of two empirical studies that evaluate these claims, and suggest that both arguments have a modest effect on people’s attitudes and behaviors regarding global poverty. In a replication of the second study, the negative duty argument, in particular, had a statistically significant effect on donations. We discuss the theoretical and practical significance of these results and suggest directions for further research on the role that philosophical arguments can play in engendering concern and action on pressing moral problems.

1 Introduction

From early in the history of the discipline, philosophers have often had the ambition of challenging and influencing people’s beliefs and, sometimes, impacting their behaviors

✉ Matthew Lindauer
matthew.lindauer@brooklyn.cuny.edu

Luke Buckland
luke.buckland@gmail.com

David Rodríguez-Arias
dra@ugr.es

Carissa Véliz
carissa.veliz@philosophy.ox.ac.uk

Extended author information available on the last page of the article

through philosophical argumentation. Plato's dialogues tell us of how Socrates traveled around Athens posing questions to citizens in order to challenge their inherited views, leading them to realize that they knew less than they thought about a variety of concepts, and particularly the concept of justice (*Republic*, Book 1). A great deal of prominent contemporary work in moral and political philosophy examines pressing moral problems and seeks both to improve our understanding of these problems and contribute to solving them. In this paper, we focus on work on the problem of global poverty, in particular, and prominent work seeking to address this problem (Singer 1972; Pogge 2002).

However, one may be skeptical of the ability of philosophical arguments to have any significant impact with respect to this type of issue. Eric Schwitzgebel and Joshua Rust (2016) have presented the results of empirical studies suggesting that professional ethicists act no differently than professors in other disciplines. These results may be thought to speak against the power of moral arguments to motivate changes in behavior. If these arguments do not improve the behavior of the people who construct and advance them, how can we expect arguments to have positive effects on the general public?

These interesting findings, however, do not seem to have stopped practically-minded ethics in its tracks. Moral philosophers continue to advance views about what is right and wrong with the goal of convincing people and, sometimes, affecting their conduct. And it does seem that at least some work in philosophy has been successful in motivating action. For instance, the animal rights movement was heavily influenced by Peter Singer's landmark book *Animal Liberation* (1975), and the field of animal ethics has inspired policy changes and movements such as the "Great Ape Project," which defends the rights of non-human great apes. The Effective Altruism movement was also heavily influenced by Singer's work on global poverty, some of which we focus on here (Singer 1972).

In the literature on global poverty, philosophers have used two main types of philosophical argument to support the claim that wealthy people have moral duties to the global poor: "positive duty" arguments and "negative duty" arguments. Peter Singer's positive duty argument (Singer 1972) and Thomas Pogge's negative duty argument¹ are among the most prominent examples of them. Positive duties are duties *to do* something, or duties to *assist* others. Singer argues that: "if it is in our power to prevent something bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it." (1972, p. 231). This general principle of *beneficence* exhorts us to aid others, regardless of whether or not we are causally responsible for their misfortunes. Negative duties, by contrast, are duties *not to do* something, e.g. not to harm others. While positive duties are commonly understood as duties of *assistance*, negative duties are understood as duties of *noninterference*. Pogge (2002) argues that, through the ongoing imposition of an unjust institutional order, the affluent have violated and continue to violate a negative duty not to harm the global poor. Importantly, the obligations resulting from these violations are duties of noninterference (i.e. negative duties), even though they may impose actions on us (e.g. compensation) (Cruft 2005). Despite taking the form of actions, compensations to the worst-off cannot be viewed as duties of assistance (positive duties), as they stem from

¹ Pogge (2002). See also Jaggard (2001).

our contributions to wrongs perpetrated against them. For Pogge, we have negative duties to the global poor because global poverty is, at least in part, our fault.² Both philosophers agree that people should donate money to organizations working to alleviate poverty, but they invoke different kinds of duties to argue for this conclusion.

Although there has been some empirical work in philosophy about how people think about giving money to relieve poverty (e.g. Hassoun et al. 2016), including on the effects of philosophical arguments (Lindauer et al. 2020), there is as yet no published evidence on which type of argument is more effective in motivating people to do so.³ The two types of arguments differ in important ways that, as noted above, many have thought could affect their relative motivational strength. On the one hand, negative duties are often regarded as more stringent than positive duties, since we intuitively seem to have weightier moral reasons to avoid harming others than we do to help them (Pogge 2002). Focusing on this difference, some have expected negative duty arguments to possess greater motivational force than positive duty arguments (Lichtenberg 2010; Lawford-Smith 2012). On the other hand, negative duty arguments rely on the premise that the affluent actively harm the global poor, a charge that may incline individuals in developed nations to defensively reject the argument as too strong and condemnatory. Positive duty arguments, by contrast, do not implicate members of developed nations in the production of global poverty, making the arguments less likely to be defensively rejected (Lichtenberg 2010). While the perceived stringency of positive duty arguments may be lesser than that of negative duty arguments—people may feel less bound to alleviate suffering that they did not cause—the lesser feeling of culpability or blame associated with positive duties may make arguments involving them more likely to produce pro-social behaviors. Judith Lichtenberg has called this the “paradoxical” aspect of the global poverty debate.⁴ However, the widely held assumption that negative duty arguments will be more motivating than positive duty arguments when accepted because they are more stringent can be tested empirically (Lawford-Smith 2012). Of course, empirical studies cannot determine the *actual* stringency of a moral duty – this is a conceptual issue. But studies can test the *perceived* stringency of moral duties, including duties to the global poor. Similarly, empirical research cannot tell us whether people *should* find positive or negative duty arguments more convincing, but it is an empirical question whether they will in fact respond to these arguments in the way that Lichtenberg has suggested, tending to more readily accept positive duty arguments.

As yet, very little attention has been paid to empirical facts about how we can be motivated to alleviate global poverty,⁵ or to the influence, if any, that philosophical arguments may have on our beliefs and behaviors regarding poverty alleviation.⁶ This

² Positive and negative duties are, of course, not incompatible. Neither Singer nor Pogge deny the existence or importance of either type of duty. However, as a consequentialist Singer is committed to the rejection of the view that negative duties are more stringent than positive duties.

³ There is also a good deal of related empirical work on charitable giving behavior, even if it does not focus on philosophical arguments. For instance, Erlandsson et al. (2018) have found that negative charity appeals (emphasizing the bad consequences of not helping) are less effective than positive ones (emphasizing the good consequences of helping), and this result may have implications for different types of philosophical arguments.

⁴ Lichtenberg 2010, p. 577–8.

⁵ Some noted exceptions include Small and Loewenstein (2003) Small et al. (2007).

⁶ Lindauer et al. 2020. On the effect of reasons on moral behavior more broadly, see Horne et al. 2015, May (2018) and Stanley et al. 2018.

is an important omission, both because prominent arguments have been thought to be better and worse than one another in respects that are relevant to their influence on us, and because much recent work on global poverty has centered on the question of feasibility, which we now turn to.

2 The Question of Feasibility

We can easily distinguish between the soundness of a moral argument (i.e. whether it is valid and contains only true premises) and its persuasiveness and motivational strength (i.e. its capacity to compel assent and affect behavior). Holly Lawford-Smith (2012) has dubbed the first type of concern “the question of truth”, and the second, “the question of feasibility” (p. 2). A feasible theory of global justice is one that will be likely to reach or contribute to (what it stipulates as) desired outcomes, such as the alleviation of global poverty. If this is right, then a feasible theory of global justice ought to take into account what will actually motivate individuals to bring these desired outcomes about. In these terms, the studies presented here focus on the question of feasibility. They examine the degree to which philosophical arguments are effective in (1) persuading people that they have duties towards the global poor, and (2) changing people’s willingness to donate money to charities that address global poverty. These two aspects of the question of feasibility are important if philosophers aspire to provide arguments that are not only sound but also capable of changing beliefs and, sometimes, motivating certain kinds of actions. In the field of global justice, in particular, issues of feasibility are important for additional reasons. Having empirically verifiable evidence about which arguments are most motivationally effective may also be of use to policymakers and NGOs as they develop campaigns that aim to alleviate global poverty.

In this paper we present the results of a set of experiments designed to empirically test the perceived plausibility and actual motivational strength of the two arguments discussed above.⁷ These experiments compare the relative effectiveness of Singer’s positive duty argument and Pogge’s negative duty argument in convincing people living in affluent countries that they have moral duties to the global poor and to meet these duties through charitable giving. Study 1 explores the impact of these arguments on people’s beliefs regarding their duties to the global poor, and bears on the perceived plausibility of each argument for non-philosophers. Study 2 and the replication study accompanying it investigate whether these arguments can affect people’s charitable giving behavior, and bear on whether the arguments differ in their motivational force.

⁷ The two types of studies described in this paper differ in various ways because the authors discovered that they were working simultaneously on the same research questions pertaining to these arguments. Each type of study contributes to our understanding of these same questions, and for this reason we decided to team up and write the present paper together.

3 Study 1: Effects of Positive and Negative Duty Arguments on Beliefs Regarding Global Poverty

In our first study, we sought to examine whether Singer's positive duty argument and Pogge's negative duty argument could be used to change people's beliefs regarding their duties to the global poor. While addressing global poverty is primarily about encouraging the right kinds of actions of states and individuals, such as reforming policies and increasing donations, philosophers typically aim to inspire moral action through rational persuasion. In contrast with other forms of persuasion, rational persuasion works through belief formation and requires that people consider the reasons that favor the conclusions that they are being persuaded to accept. Further, it is consistent with changes in behaviors being most important that changing people's beliefs will be one of the most important ways of improving their behavior over time, insofar as beliefs are reasonably stable propositional attitudes that people call on in deciding how to act. For these reasons, examining whether these arguments can affect people's beliefs about global poverty is important and practically relevant.

To examine the effects of exposure to these arguments on people's beliefs, we ran a study where undergraduate students from the United States and Spain read these arguments and tested whether their beliefs about their duties to address global poverty were affected. Participants were randomly assigned to one of three conditions, the positive duty argument condition, the negative duty argument condition, or a non-moral control condition. In all three conditions, participants first received background information about the distribution of global wealth, the effects of poverty on infant mortality, and the relatively low expense of mitigating these effects. These facts were included to ensure that any differences between the argument conditions and the control condition would not be attributable to the fact that participants in the control condition had not been thinking about the problem of global poverty. After receiving this information, participants in each condition read a vignette and argument pair. The vignettes in each condition described a scenario that matched the more general claims made in the argument, where the argument was presented as a set of numbered premises and conclusions. The positive duty vignette and argument pair read as follows:

Positive Duty Vignette

Imagine that you have inherited a successful company from a distant relative, which provided for you in your youth. On your way home from work at the company, you pass by a park with several ponds. One day while walking past one of these ponds, you see a young girl that has fallen in. She appears to be drowning. You look around—but no one else is there. If you go into the pond, you'll ruin your favorite pair of expensive shoes. You have played no role in causing her misfortune. Legal considerations aside, would you sacrifice your shoes to help save her? Now imagine that the girl is instead dying in a far-off country because her parents can't afford a simple treatment for diarrhea—a treatment that you can provide for the cost of an expensive pair of shoes. Would you save this girl? Should you save her? It's not just a thought experiment. There are children dying right now because of poverty. These are needless deaths, preventable deaths.

Positive Duty Argument

1. A great deal of extreme poverty exists, which involves suffering and death from hunger, lack of shelter, and medical care. Roughly a third of human deaths (some 50,000 daily) are due to poverty-related causes.
2. If you can prevent something bad from happening, without sacrificing anything nearly as important, you ought to do so and it is wrong not to do so.
3. By donating money to trustworthy and effective aid agencies that combat poverty, you can help prevent suffering and death from lack of food, shelter, and medical care, without sacrificing anything nearly as important.
4. Countries in the world are increasingly interdependent: you can improve the lives of people thousands of miles away with little effort.
5. Your geographical distance from poverty does not lessen your duty to help. Factors like distance and citizenship do not lessen your moral duty.
6. The fact that a great many people are in the same position as you with respect to poverty does not lessen your duty to help. Regardless of whether you are the only person who can help or whether there are millions of people who could help, this does not lessen your moral duty.

Conclusion

Therefore, you have a moral duty to donate money to trustworthy and effective aid agencies that combat poverty, and it is morally wrong not to do so.

Participants in the negative duty condition read the following vignette and argument:

Negative Duty Vignette

Imagine that you have inherited a successful company from a distant relative, which provided for you in your youth. One day, the CEO of the company retires and you make an unexpected discovery: for years the company has been knowingly building a great many defective bridges in order to increase its profits. The day after making this discovery, you walk by a pond after work and see a young girl that has fallen in as one of your company's defective bridges broke. She appears to be drowning. You look around—but no one else is there. If you go into the pond, you'll ruin your favorite pair of expensive shoes. You have unknowingly played a role in causing her misfortune. Legal considerations aside, would you sacrifice your shoes to help save her? Now imagine that the girl is instead dying in a far-off country because her parents can't afford a simple treatment for diarrhea—a treatment that you can provide for the cost of an expensive pair of shoes. The poverty this girl faces is largely the result of decades of exploitation and theft by Western governments and companies—a situation from which you have benefited by having been born in a privileged position. Would you save this girl? Should you save her? It's not just a thought experiment. There are children dying right now because of poverty that is the result of this ongoing injustice. These are needless deaths, preventable deaths.

Negative Duty Argument

1. A great deal of extreme poverty exists, which involves suffering and death from hunger, lack of shelter, and medical care. Roughly a third of human deaths (some 50,000 daily) are due to poverty-related causes.
2. Poverty is largely the result of a global system that has been constructed by the wealthy for their own benefit. This system has been imposed on the impoverished through a history of theft, violence, and exploitation.
3. You live in better circumstances than most of the world's population. You have participated in the system that has produced and maintained poverty and you have benefited from it. As a result, you have (indirectly and unknowingly) harmed the worse-off.
4. If you can avoid harming others, you have a moral duty to avoid harming them, and it is wrong not to do so. If you have harmed others, contributed to their harm, or benefited from those harms, you have a moral duty to compensate those who have been harmed, and it is wrong not to do so.
5. By donating money to trustworthy and effective aid agencies that combat poverty, you can help prevent suffering and death and compensate the poor for the harms of poverty. You can do this with very little sacrifice.

Conclusion

Therefore, you have a moral duty to donate money to trustworthy and effective aid agencies that combat poverty, and it is morally wrong not to do so.

And participants in the control condition read the following vignette and argument about the positive health effects of exercising outdoors:

Control Condition Vignette

Imagine that you have inherited a successful company from a distant relative, which provided for you in your youth. On your way home from work at the company, you pass by a park with several ponds. One day while walking past one of these ponds after work, you meet a professional athlete who explains to you the benefits of exercising outdoors, rather than indoors. She explains that exercising outdoors has the additional benefit of helping to prevent depression, and is more effective in reducing stress and anger than indoor exercising. Exercising in natural environments is associated with greater feelings of revitalization, increased energy and positive relations, together with decreases in negative emotions. People who exercise outdoors are also more likely to maintain regular exercise habits.

Given that you have easy access to outdoor spaces in which to exercise, would you change your exercise habits? It's not just a thought experiment. Millions of people in the world suffer from depression, and stress is one of the leading causes of anxiety and disease. These problems are avoidable.

Control Condition Argument

1. A great many of the health and well-being problems of the modern world have their cause in lifestyles that are increasingly sedentary (inactive) and lived indoors.
2. If there are cheap, easy and effective ways of staying healthy, you should use them to improve your health and the overall quality of your life.
3. You should try to prevent health problems and increase your well-being not only for your own benefit, but also for the benefit of the people who are close to you. This will also serve to help save society an unnecessary medical burden.
4. By frequently exercising in natural outdoor spaces, you can reduce the likelihood of suffering from health problems such as stress and depression and can increase your overall well-being.
5. Exercise in the outdoors with friends is a great way of maintaining positive social connections, which brings its own long term benefits to overall health and well-being.

Conclusion

Therefore, you should frequently exercise in natural outdoor spaces to enjoy the many health benefits it brings.

After reading the vignette and argument pair that they were assigned to, participants in all conditions were asked to rate their levels of agreement with a set of statements regarding their moral duties and beliefs about global poverty on a seven-point scale (0–6, with 0 labeled “Strongly disagree”, 3 labeled “Neither”, and 6 labeled “Strongly agree”). The two statements that we were most interested in were “People have a moral duty to donate money to trustworthy and effective aid agencies that combat poverty” and “It is morally wrong NOT to donate money to trustworthy and effective aid agencies that combat poverty, if a person is able to.” Both of these statements were included because of the way participants might interpret them. The first statement, regarding the moral duty to donate, might be viewed by some participants as too demanding because it doesn’t specify whether this duty applies even to people with very little money. We thus included the second statement, which specifies that a person must be able to donate for this duty to kick in, to help ensure that the measure was sensitive to the beliefs of participants who might be disinclined to endorse an unrestricted claim about our duty to donate, yet nonetheless believe that the relatively well-off who donate nothing have violated a moral duty.

We hypothesized that, contrary to the claim made by Lichtenberg (2010) that the positive duty argument would be easier to accept than the negative duty argument, both arguments would perform similarly to one another. Although the study does not track participant belief change directly (e.g. using a pretest-posttest design), we predicted that there would not be a significant difference between the effects that each argument had on participants’ levels of agreement with the prompts regarding duties to the poor.

3.1 Methods

3.1.1 Participants

Participants in Study 1 were 217 Spanish and 224 American undergraduate students who were recruited during their undergraduate classes. Students were an average of 22.1 years of age and the sample was roughly gender balanced, though there were slightly more women than men (women were 58% of the Spanish participants and 62% of the American participants). On average, the students claimed to have donated 2% of their income to aid agencies working on global poverty over the past 12 months (2.33% for Spanish and 2.2% for American participants).

3.1.2 Materials and Procedure

Participants were randomly assigned to one of the three conditions described above (positive duty argument condition, negative duty argument condition, control condition). The materials were initially developed in English, and members of the research team then translated them into Spanish for the Spanish group. A competent bilingual

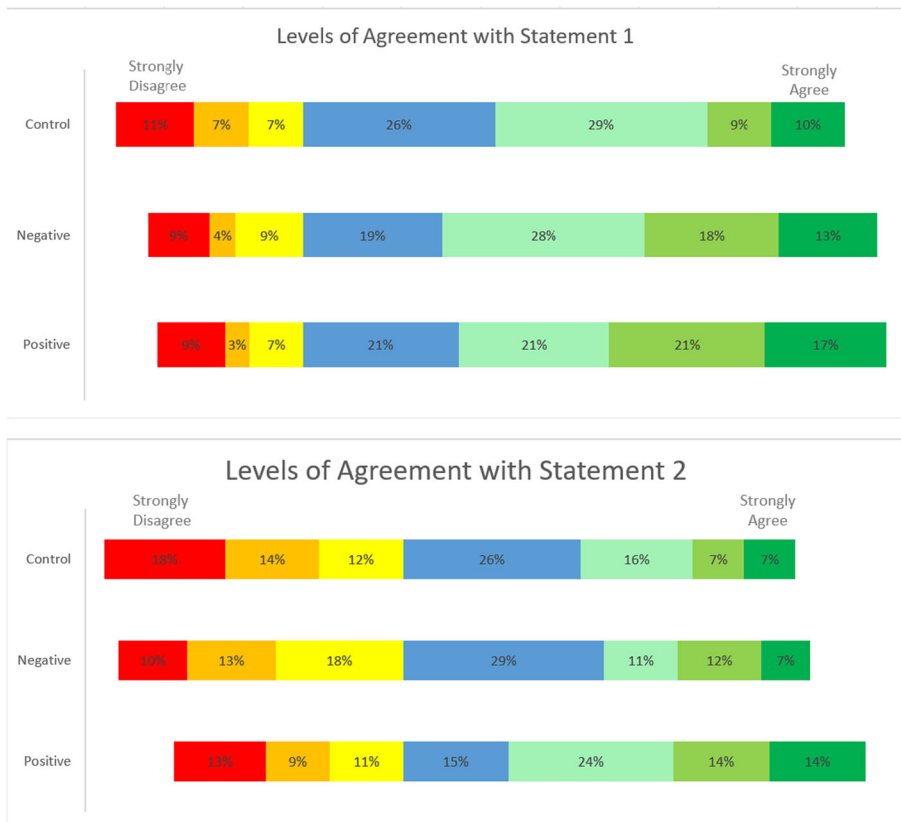


Fig. 1 Mean Levels of Agreement in Response to Statement 1 and 2

speaker not affiliated with the research team then performed a Spanish back-translation into English, which the team judged as a match of the original.

3.2 Results

Figure 1 represents participants' mean levels of agreement with the two test statements (Statement 1: "People have a moral duty to donate money to trustworthy and effective aid agencies that combat poverty" and Statement 2: "It is morally wrong NOT to donate money to trustworthy and effective aid agencies that combat poverty, if a person is able to") in response to each condition, with the vertical axis representing participants' expressed level of agreement on the seven-point Likert scale.

A 2 (Spanish, US) X 3 (positive duty argument, negative duty argument, control argument) factorial ANOVA was conducted on participants' expressed levels of agreement with each of two test statements, one on moral duty and the other moral wrongness. For the first statement, which asserted that people have a moral duty to donate to help alleviate global poverty, there was a significant main effect of argument type, $F(2, 434) = 3.447$, $p = .033$, $\eta^2 = .016$. There was no significant main effect of country and no significant interaction effect. Given the normal distribution of the data and homogeneity of group variances, Tukey's HSD post hoc test was used to compare the constituent groups; it revealed a significant difference between the positive argument condition ($M = 3.72$, $SD = 1.77$) and the control condition ($M = 3.22$, $SD = 1.69$), $p = .035$. The negative condition ($M = 3.63$, $SD = 1.69$) was not significantly different from either the positive condition, $p = .90$, or the control condition, $p = .11$.

For the second statement, which asserted that it would be morally wrong for someone not to donate to help alleviate global poverty, there was also a significant main effect of argument type, $F(2, 434) = 4.791$, $p = .009$, $\eta^2 = .022$. There was a significant main effect of country, Spain being higher, $F(1, 434) = 6.570$, $p = .011$, $\eta^2 = .015$, but no significant interaction effect. Post hoc tests using Tukey's HSD revealed a significant difference between the positive argument condition ($M = 3.25$, $SD = 1.92$) and the control condition ($MD = 2.61$, $SD = 1.80$), $p = .007$. The negative condition ($M = 2.82$, $SD = 1.68$) did not significantly differ from either the positive condition ($p = .11$) or the control condition ($p = .59$).

The analysis thus revealed a significant difference between the positive argument type and control conditions for each of the two questions. Beyond this finding, there were no other significant differences between any of the other groups assigned to the positive and negative argument types, and there were no significant differences between the participants grouped by the demographic variables.

3.3 Discussion

The results show that only participants' responses to the positive duty argument differed significantly from their responses to the non-moral control argument. The Singer-style arguments were thus more effective in changing participants' judgments about the moral value of charitable giving than were the non-moral control argument. This result gives us reason to doubt that philosophical arguments are wholly ineffective in changing people's judgments.

Despite this finding, participants' responses to the positive and negative duty arguments did not differ significantly, nor was there a significant difference between the negative duty and control arguments. These findings therefore do not support the claim that the philosophical arguments under consideration have significantly different effects on participants' subsequent judgments about their perceived moral duties, at least in the format and context in which they were presented and at the level of sensitivity of this study design. The lack of a significant difference between the positive and negative duty groups also fails to corroborate Lichtenberg's hypothesis that negative duty arguments are more likely to be defensively rejected than positive duty arguments.

4 Study 2: Effects of the Arguments on Charitable Giving Behavior

Study 1 examined the effects of reading the positive duty and negative duty arguments on people's judgments regarding global poverty, but left aside the question whether these arguments could influence their behavior. In Study 2, we sought to test whether reading these arguments could have positive effects on participants' charitable giving behavior and whether the effects of the arguments would differ from one another. We hypothesized that while both arguments would be effective in raising levels of charitable giving compared to no argument at all, neither would be more effective than the other. If the results of the study supported these hypotheses, we would have empirical evidence against the view held by Lichtenberg (2010) and Lawford-Smith (2012) that the negative duty argument is likely to be more motivationally effective than the positive duty argument. Similarly to Study 1, Study 2 and the replication study we present along with it do not aim to track changes in the beliefs of participants that may be the result of exposure to the arguments. Instead, these studies aim to detect whether there are differences in the donation behaviors between the groups exposed to the different arguments.

4.1 Methods

4.1.1 Participants

The participants in Study 2 were 90 Amazon Mechanical Turk ("MTurk") users. We limited our participant pool in the study to MTurk users from the United States with a 98% or better approval rate for their work.

4.1.2 Materials and Procedure

The study randomly assigned participants to one of three conditions: the positive duty argument condition, the negative duty argument condition, or the control condition in which no argument was given. The following versions of the arguments were used:

Singer's Positive Duty Argument

- Premise 1. We ought to give up what has no moral significance in order to prevent something very bad from happening.
- Premise 2. Giving up a luxury or frill, such as a new pair of designer shoes, is not a morally significant loss.
- Premise 3. By giving up a luxury or frill, such as a new pair of designer shoes, and donating the savings to an effective aid agency, we can save the life of a child who would otherwise die from poverty-related causes.

Conclusion

We ought to give up at least some luxuries and frills and donate the savings to aid agencies that effectively combat diseases like malaria.

Pogge's Negative Duty Argument

- Premise 1. We ought not avoidably act in ways that foreseeably harm others severely.
- Premise 2. Billions suffer today from severe poverty, which persists due to rising global inequality. Roughly a third of human deaths (some 50,000 daily) are due to poverty-related causes.
- Premise 3. The poorer half of humanity has been reduced to less than 3% of global income. The massive poverty associated with this enormous and still rising inequality is a foreseeable effect of global economic rules shaped by and for the affluent; these rules could clearly be designed to avoid much of that severe poverty.
- Premise 4. Our government plays a leading role in shaping and imposing the existing global economic rules, and we citizens bear a shared moral responsibility for what our government does in our name.

Conclusion

We ought to help protect the world's poor from the known effects of the present global economic rules by supporting organizations that effectively alleviate severe poverty and by pressuring our government to help make these rules less biased toward the affluent.

The versions of the arguments used in this study were also developed with input from both philosophers. If assigned to one of the argument conditions, the participant was asked to read the argument carefully and rate its plausibility. The plausibility question was phrased as follows: "How plausible do you find this argument (the premises, conclusion, and link between them)?" Participants registered their responses to this question on a seven-point Likert scale (with the labels "Not at all" at 1, "Somewhat" at 4, and "Extremely" at 7). While we recognized that this would not guarantee that participants were sensitive to the validity and soundness of the arguments, we hoped that the mention of the plausibility of the premises and conclusion and the "link between them" would suggest these criteria to participants in a quick and intuitive way. Following their assessment of the plausibility of the argument they received, or at the beginning of the study if assigned to the control condition,

participants were told that they had a 1 in 30 chance of being randomly selected to receive \$20. They were then asked how much of the \$20 they would want to pay forward to an organization working to address the poverty-related deprivations of persons in developing countries. If assigned to an argument, they were asked how much they would want to pay forward to an organization working to address the poverty-related deprivations of persons in developing countries on the basis of the argument that they read. In the control and argument conditions, participants were told that the money that they chose to pay forward would go to one of the top organizations working on global poverty as rated by GiveWell. The GiveWell website address (<http://www.givewell.org>) was included with this statement, and participants were also told that GiveWell is an independent institution dedicated to reviewing organizations' effectiveness along a number of strict, testable criteria. The information regarding GiveWell was included as an attempt to counteract skepticism that participants might have regarding the effectiveness of charitable donations.⁸

4.2 Results

Figure 2 represents the mean donation amounts that participants chose in response to each condition (positive duty argument, negative duty argument, control).

The plausibility ratings that participants gave to Singer's positive duty argument ($M = 5.17$, $SD = 1.05$) and Pogge's negative duty argument ($M = 5.13$, $SD = .90$) did not differ significantly, $t(58) = .132$, $p = .90$. The mean plausibility ratings for both arguments were above the midpoint, suggesting that participants generally regarded both arguments as plausible.

A one-way ANOVA was conducted on participants' mean donation amounts in response to the control condition ($M = 4.70$, $SD = 4.57$), positive duty argument ($M = 9.12$, $SD = 5.82$), and negative duty argument ($M = 8.63$, $SD = 6.01$).⁹ This analysis confirmed that there was a significant main effect of condition.¹⁰ Post hoc tests using Tukey's HSD revealed that while there was no significant difference between donation amounts in response to the two arguments ($p = .936$), donation amounts in response to the positive duty argument and negative duty argument were both significantly greater than donation amounts in response to the control condition ($p = .007$, $p = .019$, respectively).

⁸ In this study and the replication study that follows it, winning participants were paid their bonuses, and their remainders that they selected to donate were paid to a top-rated GiveWell charity at the time of payment.

⁹ The donation amounts given in response to the control condition were not normally distributed (Shapiro-Wilk, $W(30) = .87$, $p = .001$), largely because subjects in that condition tended to offer to give very small amounts. While the one-way ANOVA is robust to violations of normality at large enough sample sizes, it's worth noting that we also ran a Kruskal-Wallis H test on the data, a nonparametric test that can be used to analyze data that isn't normally distributed. Because the shape of the distribution of responses to the control condition differs from the shapes of the distributions of responses to the argument conditions, we use the Kruskal-Wallis H test to compare the mean ranks of the groups. This test also showed that there was a statistically significant difference between donation amounts in response to the three conditions, $\chi^2(2) = 12.20$, $p = .002$, $\eta^2 = .14$. Post-hoc tests showed that the donations given in response to the positive and negative duty arguments were significantly greater than those given in response to the control condition ($p = .004$ and $p = .018$ respectively), with no significant difference between the argument conditions ($p = 1.000$).

¹⁰ $F(2, 87) = 5.81$, $p = .004$, $\eta^2 = .12$.




Fig. 2 Mean Donation Amounts Selected in Response to Each Condition

4.3 Discussion

The results of Study 2 suggested that both the positive and negative duty arguments could have significant effects in increasing the amount of money that participants are willing to donate to organizations working to address global poverty. In Study 2, neither argument had a greater effect than the other on the amounts chosen by participants to have paid forward.

However, through the review process for this paper, we came to realize that Study 2 had a number of shortcomings that we should improve upon in a replication study.¹¹ First, the study was underpowered, needing a larger sample size to be sure that we were detecting a genuine effect of the arguments on donations.¹² Second, because the control condition didn't contain any argument, it might be valuable to compare the argument conditions to a control condition that did contain one, such as a non-moral argument.

In a replication for Study 2, we improved upon the study's design in these ways, using a larger sample size and a non-moral argument control condition, similar to the control argument used in Study 1.¹³ If either or both of the moral arguments still had a positive effect on donations, we would have more evidence that moral arguments can produce improvements in charitable giving. Our specific hypotheses were: (1) Exposure to both positive duty and negative duty arguments for the moral duty of charitable giving will elicit higher levels of charitable giving from participants as compared to a

¹¹ We are grateful to an anonymous reviewer for this and other feedback that led us to conduct the replication study.

¹² We used GPower 3.1 software to conduct a power analysis. For an ANOVA analyzing the observations in the three conditions (positive duty, negative duty, control), we would need 206 participants to have a 90% chance of detecting a medium sized effect ($f = .25$). However, we will need a slightly larger sample to conduct our post hoc analyses examining the differences between the two argument conditions and the differences between each argument and the condition. To have a 90% chance of finding a medium sized difference between these groups ($d = .5$) we would need 86 participants per condition, with three conditions totalling to 258 participants.

¹³ The replication is registered at <https://osf.io/d29nr>.




Fig. 3 Mean Donation Amounts Selected in Response to Each Condition

non-moral control argument unrelated to the fulfillment of charitable duties; and (2) The positive and negative duty arguments will be equally effective in eliciting charitable giving from participants, and thus we will observe no statistically significant differences between the levels of charitable giving of the groups of participants exposed to each argument.

The procedure for the replication study was the same as for original Study 2, except for the addition of the following non-moral argument that was used in a control condition:

Non-Moral Argument

- Premise 1. A great many of the health and well-being problems of the modern world have their cause in lifestyles that are increasingly sedentary (inactive) and lived indoors.
- Premise 2. If there are cheap, easy and effective ways of staying healthy, you should use them to improve your health and the overall quality of your life. For both your sake and for the sake of the people who are close to you, as well as to save society an unnecessary medical burden.
- Premise 3. By frequently exercising in natural outdoor spaces, you can reduce the likelihood of suffering from health problems.

Conclusion

Therefore, you should frequently exercise in natural outdoor spaces to enjoy the many health benefits it brings.

Plausibility ratings were similar across the Positive Duty ($n = 89$, $M = 5.38$, $SD = 1.45$), Negative Duty ($n = 87$, $M = 5.45$, $SD = 1.36$), and Non-Moral Control ($n = 89$, $M = 5.38$, $SD = 1.19$) conditions and did not differ significantly from one another

($p > .05$). Figure 3 represents the mean donation amounts given in response to each argument.

Mean donations were higher for the Positive Duty argument ($M = 7.52$, $SD = 6.69$) and Negative Duty argument ($M = 8.3$, $SD = 6$) than the Non-Moral Argument ($M = 6.31$, $SD = 6.27$). This provides some evidence in favor of our first hypothesis.

Donations in response to each condition were not normally distributed,¹⁴ so we used non-parametric tests to evaluate differences in donation amounts. Mann-Whitney U tests showed that the difference between the Negative Duty condition and Non-Moral control condition was significant ($U = 3127.5$, $p = .026$), whereas the other pairwise comparisons, and notably the comparison between the two arguments conditions, were not,¹⁵ with this latter comparison providing evidence to support our second hypothesis. The lack of significance for the Positive Duty condition as compared to the Control condition, which also casts some doubt on our first hypothesis, may be due to the relatively low sample size for a donation study, even though it meets the requirements of the power analysis. We cannot reliably conclude that the Positive Duty argument did not lead to increased donations. In other research, with a larger potential payoff, it has been observed that a similar type of Positive Duty argument can significantly increase charitable giving (Lindauer et al. 2020). But we did not find a statistically significant effect of the Positive Duty argument on donations in the replication of Study 2.

As noted, in the replication study hypothesis 2 was supported, with no significant difference in the effects of the two arguments, although the negative duty argument surpassed the threshold for statistical significance while the positive duty argument did not. These results taken together provide mixed evidence regarding the empirical speculation in the global poverty literature that negative duty arguments will tend to be more motivating than positive duty arguments (Lichtenberg 2010; Lawford-Smith 2012).

Notably, Study 2 and the replication study employ a potential payoff model. There are therefore limits to the conclusions we can draw about whether philosophical arguments can affect charitable giving behavior “outside the lab” on the basis of our results. However, it is not uncommon in the literature on decision-making and donations for studies to use potential payoff models (see, e.g., Soyer and Hogarth 2011; Lindauer et al. 2020). As with other studies using these models, we cannot rule out the possibility that our results will not carry over to participants’ charitable giving using their own funds. This is a topic for further research. Still, we take our results to provide some evidence that these arguments can have motivational force, and that their motivational force may not be very different.

5 Taking Stock

In this paper we examined whether different philosophical arguments have the power to affect both people’s judgments and their behaviors in relation to global poverty. Our results suggest that they can in fact do so. They also suggest that the prominent

¹⁴ Positive Duty condition, $W(89) = .888$, $p < .001$; Negative Duty condition, $W(87) = .928$, $p < .001$; Control condition, $W(89) = .869$, $p < .001$.

¹⁵ Positive Duty vs. Control, $U = 3529.5$, $p = .205$; Positive Duty vs. Negative Duty, $U = 4168$, $p = .378$.

arguments that we tested, Singer's positive duty argument and Pogge's negative duty argument, have similar effects on these judgments and behaviors.

Of course, we must leave aside a number of questions outside the scope of the present paper. In particular, we do not examine what in the arguments, precisely, participants were responding to. Our results also raise questions about the comparative stringency of positive and negative duties that we cannot fully address here. They do suggest, however, that the common assumption that negative duties are more stringent than positive duties may not be in line with people's ordinary judgments when thinking about the issue of global poverty. If moral stringency is established in relation to ordinary people's judgments, both types of duties towards the global poor may be similar in their moral stringency, although perhaps the replication study suggests that negative duties are more stringent. Alternatively, stringency may be established in a way that is not influenced by commonsense intuitions or perceived moral weight; moral philosophers may choose to argue that the general public is simply wrong about the relative stringency of positive and negative duties regarding global poverty. Philosophical theories that defend a view of the stringency of these duties that diverges from ordinary people's judgments would, of course, benefit from having a plausible account of why people fail to get the facts about the relative stringency of these duties right. Further empirical studies would be needed to explore each of these issues.

Our results suggest that the predictions offered by Lichtenberg (2010) and Lawford-Smith (2012), inspired by the common way of looking at the relative costs and benefits of each type of argument, require further investigation. On the one hand, the common view seems to be mistaken in regarding the positive duty argument as easier for people to accept than the negative duty argument. On the other hand, the common view also appears to be mistaken in regarding the negative duty argument as more motivating than the positive duty argument, although as we saw, in the replication study negative duties achieved a statistically significant result whereas the positive duty argument did not. Further studies with larger samples will need to be conducted to examine these potential contrasts in greater depth.

We should note that a general perceived lack of control with regard to global poverty, perhaps the product of pessimism about the possibility of eradicating it, could have affected our results, evening out the effectiveness of both arguments.¹⁶ It is also possible that Pogge's argument may not be a paradigmatic instance of a negative duty argument because, as Steinhoff argues (2012), it would take a great deal of action to avoid harming the global poor, and our involvement in the production of the harm itself does not seem like an action to most people.

¹⁶ It is less likely that the perceived lack of control is having an impact on the results of Study 2 and the replication study. These studies included direct mention of the effectiveness of the charities that are appealed to, pointing subjects to GiveWell's rigorous, testable criteria for effective NGOs. Thus, it seems unlikely that skepticism about the effectiveness of aid is affecting the results.

6 Conclusion

This paper examined the relative effects of Singer's positive duty argument and Pogge's negative duty argument on judgments related to the problem of global poverty and actions that aim to address it. Our results suggest that both approaches can serve to improve people's beliefs and their behaviors in relation to the global poor. This is an important finding with both theoretical and practical implications. It suggests that philosophers are not merely fooling themselves in thinking that their arguments may produce positive changes in the world. In particular, it seems that philosophical arguments can be effective in convincing people that they have moral duties to the global poor and move them to act in ways that are consistent with these convictions.

Moral philosophy provides the theoretical insights that inspire this empirical study, while moral psychology allows for moral theories to be evaluated on empirical grounds. By bringing theoretical and empirical research together, this study shows that it is possible to test whether conjectures about the real-world impact of philosophical arguments are likely to be true, and potentially find new ways of motivating people to participate in solving pressing moral problems.

Acknowledgements We are grateful to Mark Alfano, Elizabeth Anderson, Lilian Bermejo, Pedro Francés, Tamar Gendler, Joshua Knobe, Meena Krishnamurthy, Ester Massó, Juan Antonio Montero, Melania Moscoso, Jesse Prinz, Stephen Stich, and Boris Yakubchik for helpful feedback. We would also like to thank Thomas Pogge and Peter Singer for helping us develop the versions of the arguments used in Study 2 and the replication study, Jaclyn Murray for the Spanish back-translation, and Marcus Mayorga for assistance with the statistical analyses used in the paper.

Author Contribution All authors contributed equally to this paper.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Cruft, R. 2005. Human rights and positive duties. *Ethics & International Affairs* 19: 29–37.
- Erlandsson, A., A. Nilsson, and D. Västfjäll. 2018. Attitudes and donation behavior when Reading positive and negative charity appeals. *Journal of Nonprofit & Public Sector Marketing* 30 (4): 444–474.
- Hassoun, N., E. Malikov, and N. Lubchenco. 2016. How people think about distributing aid. *Philosophical Psychology* 29 (7): 1029–1044.
- Horne, Z., D. Powell, and J. Hummel. 2015. A single counterexample leads to moral belief revision. *Cognitive Science* 39(8):1950–1964.
- Jaggar, A.M. 2001. Is globalization good for women? *Comparative Literature* 53 (4): 298–314.
- Lawford-Smith, H. 2012. The motivation question: Arguments from justice and from humanity. *British Journal of Political Science* 42 (3): 661–678.
- Lichtenberg, J. 2010. Negative duties, positive duties, and the new harms. *Ethics* 120 (3): 557–578.

- Lindauer, M., M. Mayorga, J.D. Greene, P. Slovic, D. Västfjäll, and P. Singer. 2020. Comparing the effect of rational and emotional appeals on donation behavior. *Judgment and Decision Making* 15 (3): 413–420.
- May, J. 2018. *Regard for reason in the moral mind*. Oxford: Oxford University Press.
- Pogge, T. 2002. *World poverty and human rights*. New York: Polity Press.
- Schwitzgebel, E., and J. Rust. 2016. The behavior of ethicists. In *A Companion to Experimental Philosophy*, ed. J. Sytsma and W. Buckwalter, 225–233. Hoboken: Wiley-Blackwell.
- Singer, P. 1972. Famine, affluence and morality. *Philosophy and Public Affairs* 1: 229–243.
- Singer, P. 1975. *Animal liberation: A new ethics for our treatment of animals*. New York: Harper Collins.
- Small, D.A., and G. Loewenstein. 2003. Helping a Victim or helping the victim: Altruism and Identifiability. *Journal of Risk and Uncertainty* 26: 5–16.
- Small, D.A., G. Loewenstein, and P. Slovic. 2007. Sympathy and callousness: The impact of deliberative thought on donations to identifiable and statistical victims. *Organizational Behavior and Human Decision Processes* 102: 143–153.
- Soyer, E., and R.M. Hogarth. 2011. The size and distribution of donations: Effects of numbers of recipients. *Judgment and Decision making* 6: 616–628.
- Stanley, M.L., A.M. Dougherty, B.W. Yang, P. Henne, and F. De Brigard. 2018. Reasons probably Won't change your mind: The role of reasons in revising moral decisions. *Journal of Experimental Psychology: General* 147 (7): 962–987.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Luke Buckland¹ · Matthew Lindauer² · David Rodríguez-Arias³ · Carissa Véliz⁴

¹ The Center for Leadership Ethics in Africa, University of Fort Hare, Alice, South Africa

² Department of Philosophy, Brooklyn College, City University of New York, 2900 Bedford Avenue, Brooklyn, NY 11210, USA

³ FiloLab-UGR, Philosophy I Department, Universidad de Granada, Granada, Spain

⁴ Institute for Ethics in AI, University of Oxford, Oxford, UK