

The many obstacles to effective giving

Lucius Caviola*

Stefan Schubert[†]

Jason Nemirow[‡]

Abstract

When people donate, they rarely give to the charities that do the most good per dollar. Why is this? One possibility is that they do not know how to give effectively. Another possibility is that they are not motivated to do so. Across six tasks (Studies 1a, 1b), we found support for both explanations. Among lay donors, we observed multiple misconceptions—regarding disaster relief, overhead costs, donation splitting, and the relative effectiveness of local and foreign charities—that reduced the effectiveness of their giving. Similarly, we found that they were unfamiliar with the most effective charities (Studies 2a, 2b). Debunking these misconceptions and informing people about effectiveness boosted effective donations; however, a portion of lay donors continued to give ineffectively to satisfy their personal preferences. By contrast, a sample of self-identified effective altruists gave effectively across all tasks. They exhibited none of the misconceptions that we observed among lay donors and overwhelmingly favored the most effective option in their choice set (Study 3). Taken together, our studies imply that donors need to be both informed and motivated to give effectively on a consistent basis.

Keywords: charitable giving, effective altruism, prosocial behavior, misconceptions, preferences

1 Introduction

Each year, people donate more than \$400 Billion to charity in the US alone — equivalent to 2% of American GDP. Yet, studies show that people often donate to charities that save fewer lives, or otherwise do less good than the most effective charities (Bagwell, de Las Casas, van Poortvliet & Abercrombie, 2013). Since the most effective charities are much more effective than the average charity (Caviola et al., 2020), this means that most donors do much less good with their donations than they could have done (Lomborg, 2012; Jamison et al., 2006; Tollefson, 2015). Only a relatively small number of donors, e.g., in the Effective Altruism movement (MacAskill, 2015), systematically donate to the most effective charities, in order to do the most good per dollar.

Why do most people donate ineffectively? One explanation is that they are not motivated to donate effectively (the preference-based explanation). In support of that ex-

planation, Berman, Barasch, Levine & Small (2018) show that people prefer to support causes they feel emotionally attached to even when they know that other options are more effective. This implies that the reason why people give ineffectively is not a lack of information, but rather preferences for an ineffective option. (By this we do not mean that donors have a preference for ineffectiveness *per se*, but rather that their preferred option happens to be ineffective.) If donors' personal preferences were to explain ineffective giving, then we should expect that informing people about how to donate effectively would not make them any more likely to do so.

Another explanation for why people donate ineffectively is that they do not know how to do so (the belief-based explanation, or “distorted altruist” explanation, see Berman et al., 2018). In support of this explanation, it has been shown that people give less to charities with high overhead costs because they suffer from the misconception that high overhead necessarily entails low cost-effectiveness (Baron & Szymanska, 2011; Caviola, Faulmüller, Everett, Savulescu & Kahane, 2014; Gneezy, Keenan & Gneezy, 2014). Once participants learn that overhead costs and cost-effectiveness can diverge, they give more to charities with the highest cost-effectiveness. The belief-based explanation says that the reason people donate ineffectively is that they suffer from such misconceptions, or that they simply lack knowledge. If this turned out to be the explanation for ineffective giving, we should expect that informing people of how to donate effectively would make them do so.

In this paper, we aim to test these explanations. Our hypothesis is that they both play a role. People donate ineffectively partly because they do not know how to donate effectively (they have false beliefs, i.e., misconceptions, and simply lack knowledge) and partly because they do not want

The first two authors contributed equally.

Reports of all measures, manipulations, materials, data (including exclusions), and analysis code are available for download at: <https://osf.io/42ke3/>.

We would like to thank Fabienne Sandkühler for her extensive help and comments, Gregory Lewis, Aaron Gertler, David Moss, Matti Wilks, Pablo Stafforini, David Althaus, Aron Vallinder, and Steven Pinker for their helpful suggestions, and the experts for their help in evaluating the charities.

Copyright: © 2020. The authors license this article under the terms of the Creative Commons Attribution 3.0 License.

*University of Oxford and Harvard University. Corresponding author: Department of Psychology, Harvard University, 33 Kirkland St, Cambridge, MA 02138. Email: lucius.caviola@gmail.com. ORCID 0000-0003-4302-5884

[†]University of Oxford.

[‡]Harvard University.

to donate effectively (they have personal preferences for less effective options). In line with the belief-based explanation, we hypothesize that providing people with information about how to donate effectively will make their giving more effective to a certain extent (Karlan & Wood, 2017; Metzger & Günther, 2019). However, in line with the preference-based explanation, we also hypothesize that people exposed to such information, no matter how credible or well-argued it may be, will not approach the ceiling of effective giving because of their personal preferences for less effective options.

We will now look at a number of concrete obstacles to effective giving and explain how our hypothesis applies to each of them. After searching the academic and industry literature on charity cost-effectiveness, we identified seven paradigmatic obstacles to effective giving. Several, but not all, of these obstacles were studied by Baron and Szymanska (2011). We extend their research by investigating whether these obstacles are due to a lack of knowledge, preferences for ineffective options, or both. Our list is neither exhaustive nor completely representative, but should suffice for our purposes.

1. Focusing on overhead costs instead of cost-effectiveness. As mentioned above, previous work has shown that people give less to charities with high overhead costs (Baron & Szymanska, 2011; Caviola et al., 2014; Gneezy et al., 2014). We hypothesize that this is both because they have an intrinsic aversion to high overhead *and* because they falsely believe that higher overhead necessarily entails lower cost-effectiveness. Recent research suggests that high overhead is, in fact, associated with *high* cost-effectiveness (Haynie, 2019). Overhead costs are not necessarily wasted, because charities need to hire competent staff and build infrastructure to accomplish their mission. Thus, when trying to determine the effectiveness of different charities, one needs to look directly at their cost-effectiveness, rather than at their level of overhead.

2. Giving to local instead of foreign charities. People in rich countries often help people in their local community rather than distant strangers (Baron & Szymanska, 2011). This is encapsulated in the popular slogan, “Charity begins at home.” We hypothesize that this is both because people have a preference to help those who are more socially and physically proximate (Nagel & Waldmann, 2013) and because they believe that it is more effective to help local people than distant people in poor countries (Knowles & Sullivan, 2017). The latter point is, however, incorrect: on the contrary, donors in richer countries can expect to do more good by giving to well-run charities in poor countries, compared to donating to similar charities in their own countries (GiveWell, 2012; MacAskill, 2015).

3. Giving to disaster relief instead of recurring problems. We hypothesize that people have a tendency to prioritize donating to disaster relief over the mitigation of a recurring problem. (As far as we know, we are the first to

demonstrate this effect experimentally.) Furthermore, we hypothesize that this is both because people have a preference for donating to disaster relief *and* because they believe that it is more effective. With respect to efficacy, the opposite appears to be true: donating to disaster relief is likely to be less effective than donating to charities that work on recurring problems (MacAskill, 2015). One reason for this is that disaster charities often receive a large influx of cash in a short time period due to heavy media coverage, but lack the infrastructure to distribute these funds effectively.

4. Avoiding charities that have a small chance of achieving a great impact. We hypothesize that people have a tendency to shy away from charities that have a small chance of having a huge impact. As far as we know, this has not yet been explicitly tested in a charitable giving context. We further hypothesize that this is because people are risk averse *and* because they genuinely believe that the “safer” option is more effective. In other words, they have a preference for a known option over an option with an uncertain outcome (e.g., because they want to be sure that their money will not go to waste), even if the latter has a higher expected value (Brock, Lange & Ozbay, 2013) *and* they believe that the safer option carries a higher expected value, even when it does not.

5. Splitting donations instead of giving everything to the most effective charity. When presented with two charities, one of which is more effective than the other, people have a tendency to split their donation across the two charities (Baron & Szymanska, 2011). We hypothesize that this is due both to a preference for splitting, perhaps driven by considerations of fairness (Sharps & Schroeder, 2019), *and* to the belief that diversifying donations is more effective than concentrating them into one charity. However, when one of the charities is more effective than the other, giving the whole donation to the most effective charity typically yields the highest impact, since the marginal impact of donating to the more effective charity stays higher even if it receives the full sum. This is partly because the difference between the most effective charities and the average charity is larger than most people believe: e.g., a factor of about one hundred in the global poverty domain, according to experts (Caviola et al., 2020).

6. Focusing on identifiable victims instead of the greatest number. People tend to prioritize donating to a single identifiable victim over larger numbers of unidentifiable victims (Slovic, 2010; Small, Loewenstein & Slovic, 2007). Unlike the other cases, we hypothesize that this is solely because people have a preference to help identifiable victims. We do not think that they fail to see that helping the unidentifiable victims is more effective. If correct, this would show that even though preferences for ineffective options and false beliefs jointly cause people to donate ineffectively in most cases, that is not always the case. On this hypothesis, there

are cases where preferences alone can explain ineffective giving.

7. Unfamiliarity with the effective charities. In the real world, people rarely select the most effective charities from the vast global pool of such organizations. We hypothesize that this is both because people have strong preferences for specific charities — which they would choose even if they turned out to be ineffective — *and* because they do not know which charities the evaluation experts, such as *GiveWell*, have judged to be highly cost-effective (e.g., charities fighting malaria and neglected tropical diseases).

1.1 The present research

To test our hypotheses, we investigated responses to the seven experimental tasks described above from laypeople (Studies 1a and 1b (cases 1–6) and Studies 2a and 2b (case 7)) and self-identified effective altruists (Study 3). Our aim was to explore the extent to which beliefs and preferences explain ineffective giving.

2 Study 1a: Six donation tasks

In Study 1a we used a range of giving tasks to test our hypothesis that beliefs and preferences typically cause ineffective giving together across a range of tasks.

2.1 Method

Participants. We recruited 421 US participants from MTurk. We excluded 17 because they failed two or more attention checks out of four (see materials on OSF). The final sample consisted of 404 participants (189 females, age $M = 42.80$, $SD = 12.70$). Our sample size was determined by a power analysis which suggested that we would need at least 351 participants to detect an effect size of $f = 0.15$, relying on an alpha of 0.05, power of 0.8, and four groups. We aimed to recruit 400 participants to account for exclusions.

Procedure. We gave participants six separate effective giving tasks, based on the first six cases discussed above, presented in randomized order (randomized for each participant). All tasks included descriptions of two donation options (e.g., Charity A and B), one of which is more effective. For each task, participants responded on a 7-point scale with anchors at *Definitely Option A (1)*, *Unsure (4)*, and *Definitely Option B (7)*.

More specifically, the six tasks featured the following options (see online materials for details): 1) a charity with low overhead costs and medium cost-effectiveness vs. a charity with high overhead costs and high cost-effectiveness (*Overhead*), 2) a local health charity vs. a foreign health charity (*Local*), 3) a disaster relief charity vs. a charity focusing

on recurring health issues in the developing world (*Disaster relief*), 4) a charity saving lives with certainty vs. a charity saving more lives in expectation but with lower certainty (*Risk aversion*), 5) an option to split donations across an effective and less effective charity vs. an option to concentrate donations in the effective charity (*Splitting*), 6) a charity focused on helping a single identifiable victim vs. a charity that helps multiple unidentifiable victims (*Identifiable victim*).

The study had a 2 (information: yes vs. no) x 2 (question type: personal-donation vs. effectiveness-belief) between-participants design. Participants in the information conditions received additional information that debunked a common misconception, and explained which donation option is likely the more effective one. For example, we informed participants in the information conditions that donating to alleviate recurring health problems in the developing world is generally more effective than donating to disaster relief. This manipulation allowed us to infer the extent to which false beliefs about charities cause ineffective giving. If false beliefs contribute to ineffective giving, then we should observe participants in the information conditions choosing the effective option more often than participants in the no-information condition.

In the personal-donation condition, participants were asked to choose the donation option they would personally donate to. Participants in the effectiveness-belief condition were asked to select the donation option they believe is the most effective. This manipulation allowed us to infer the extent to which subjective preferences for ineffective options cause ineffective giving. If personal preferences contribute to ineffective giving, then we should observe participants in the effectiveness-belief condition choosing the effective option more often than participants in the personal-donation condition.

At the end of the study, participants indicated, on a 7-point Likert scale, to what extent they want their donations to “achieve the greatest amount of good for the largest number of people possible”. Finally they responded to demographic questions.

2.2 Results and discussion

We ran six two-way ANOVAs for each of the six tasks (Table 2). Across all tasks, we found significant main effects of the information factor and the question type factor (see Figure 1 and Tables 1–2).

First, in all tasks, participants were more likely to choose the effective donation option if they were given additional information that debunked the specific misconception in question. This was the case both in the conditions where participants were asked which charity they would personally donate to and in the conditions where participants were asked which




FIGURE 1: Note that 1 stands for definitely choosing the less effective option, 4 for being unsure which option to choose, and 7 for definitely choosing the effective option. In Study 1a, MTurk participants were more likely to choose the effective donation option when asked to select the option they believed to be more effective (Belief) than when asked to select the option they would donate to personally (Donate). Participants were also more likely to choose the effective option when they were given further information that debunked a specific misconception about the relative effectiveness of the two options. In Study 3, effective altruists chose the effective donation option in all tasks regardless of whether they were informed which option is the most effective, and regardless of whether they were asked to identify the most effective option, or to choose the option they would personally donate to.

charity they believe is the most effective. Assuming that participants stated their true beliefs, this suggests that, in the no-information condition, participants chose the ineffective option partly because they had false beliefs about the relative effectiveness of the two charities; beliefs which the information manipulation corrected.

Second, in all tasks and across both information conditions, participants who were asked to select the option they believed to be the most effective were more likely to choose the effective option than participants who were asked which option they would donate to themselves. This suggests that participants in the personal-donation condition selected charities that they believed to be less effective, because they had personal preferences for those options.

Third, while we did not find interactions in most tasks, there were significant interaction effects in the identifiable victim and disaster relief tasks. In those two tasks, informing participants which option is more effective made a stronger difference when participants were asked which donation option they would personally choose, than when asked which donation option they believe is the most effective. Note, though, that it is unclear how robust these two interactions are and whether they are interpretable.

We observed similar patterns across the six tasks. In most tasks, the information manipulation made a greater difference than the question-type manipulation. The two exceptions were the local (vs. foreign) charity task and the identifiable victim task, where the question-type manipulation made a greater difference. The fact that the question-type manipulation made such a large difference in those tasks suggests that participants had strong preferences to help the identifiable victim and support the local charity — preferences which persisted even when participants learned that these options are ineffective. In fact, our results suggest that most people already knew that helping the identifiable victim is ineffective before providing them with the information (though not everyone; the information manipulation still made a significant difference between the two personal-donation conditions). The information manipulation made the biggest difference in the disaster and overhead tasks, suggesting that false beliefs were particularly strong causes of ineffective giving in those tasks.

Participants overwhelmingly stated that, when they donate to charity, they want to achieve the greatest amount of good for the largest possible number of people ($M = 6.06$; $SD = 1.00$; on a 7-point Likert scale), even though they did

TABLE 1: Means and standard deviations for each task and condition. (Study 1a).

Task	Personal-donation		Effectiveness-belief	
	No Information	Information	No Information	Information
Overhead	2.03 (1.37)	4.80 (2.04)	2.31 (1.50)	5.65 (1.60)
Local	2.33 (1.62)	3.88 (2.38)	4.33 (2.00)	5.66 (1.78)
Disaster relief	2.75 (1.68)	5.38 (1.75)	3.90 (1.94)	5.81 (1.61)
Risk aversion	2.23 (1.77)	4.07 (2.10)	2.84 (2.16)	5.12 (2.02)
Splitting	3.22 (2.48)	4.83 (2.14)	4.23 (2.65)	5.75 (1.93)
Identifiable victim	3.91 (2.03)	5.17 (1.98)	5.97 (1.65)	6.07 (1.61)

TABLE 2: ANOVAs with donation choice (from definitely effective to definitely not effective) as the dependent variable and question type (personal-donation vs. effectiveness-belief), information, and their interaction as independent variables. (Study 1a).

Task	Question type		Information		Interaction	
	F	η_p^2	F	η_p^2	F	η_p^2
Overhead	12.53**	.029	348.50**	.466	2.98	.007
Local	94.07**	.189	54.31**	.120	0.34	< .001
Disaster relief	21.25**	.049	170.03**	.298	4.16*	.010
Risk aversion	17.57**	.041	105.24**	.208	1.18	.003
Splitting	17.94**	.042	46.45**	.104	0.04	< .001
Identifiable victim	66.47**	.142	13.95**	.034	10.18**	.248

The df for all scenarios were (1,400). * $p < .05$, ** $p < .001$.

not always choose the most effective option in practice. We aggregated the donation choices across all six tasks to form a single average effective donation choice score per participant ($\alpha = .74$). We then computed a correlation between the desire to do the most good per donation and the aggregated effective donation choice score in each of the four conditions. The more a participant wanted to do the most good per donation, the more ineffectively they donated in the no-information condition, $r = -.19$, $p = .05$, and the more effectively they donated in the information condition, $r = .52$, $p < .001$. The desire to do the most good per donation did not correlate with beliefs about the relative effectiveness of the two options in the no-information condition, $r = -.01$, $p = .92$, but it did correlate with believing the information they were given about which charity is more effective in the information condition, $r = .26$, $p = .008$. These findings show that the desire to do the most good increases the chance of donating to the most effective charity only when one knows which option is in fact the most effective one.

In sum, this study confirmed our hypothesis that both misconceptions about charities' relative effectiveness and preferences for ineffective options cause ineffective giving.

3 Study 1b

In Study 1b we aimed to replicate the findings of Study 1a with a simpler design. Whereas Study 1a included four conditions, Study 1b only featured two conditions: one where possible misconceptions regarding which option is more effective were debunked (information condition), and one where there was no such debunking (no-information condition). Unlike in Study 1, each participant was asked both where they would donate personally and which option they believe is more effective. This design allowed us to study why individuals who were given information about which option is likely to be more effective still chose the less effective one.

We also made some slight changes to the vignettes and debunking texts, to ensure they were not too persuasive, and defined the concept of "cost-effectiveness" more clearly.

3.1 Method

Participants. We recruited 209 US participants from MTurk. We excluded 10 because they failed two or more

TABLE 3: Donations

	No Information	Information	t	df	Cohen's d
Overhead	3.20 (2.20)	4.48 (2.09)	4.20***	195.95	.59
Local	2.77 (1.79)	3.56 (2.24)	2.74**	177.62	.39
Disaster relief	3.22 (1.70)	5.36 (1.64)	9.04***	195.98	1.28
Risk aversion	3.14 (2.34)	4.31 (2.29)	3.55***	195.46	.50
Splitting	3.48 (2.53)	4.84 (2.25)	4.02***	197.0	.57
Identifiable victim	4.52 (1.93)	5.3 (1.76)	2.95**	190.29	.42

Note. * $p < .05$, ** $p < .01$, *** $p < .001$

attention checks (out of four). The final sample consisted of 199 participants (73 females, age $M = 40.84$, $SD = 12.45$). Our sample size was determined by a power analysis which suggested that we would need at least 200 participants to detect an effect size of $d = .4$, relying on an alpha of 0.05, power of 0.8, and four groups. We aimed to recruit 210 participants to account for exclusions.

Procedure. We gave participants six separate effective giving tasks, which were slightly altered versions of those in Study 1a, presented in randomized order. In the information condition, participants read information from charity experts that indicated which option would save more lives, along with a brief explanation of the logic behind the expert judgements. These information texts were slightly altered versions of those used in Study 1a, written in a less persuasive style.

In contrast to Study 1a, all participants responded to the same questions. First, they indicated which option they would personally donate to on a 7-point scale with anchors at *Definitely Option A (1)*, *Unsure (4)*, and *Definitely Option B (7)*. Second, they indicated which option they believed would save more lives on the same 7-point scale. Third, they indicated which option they thought had greater overall positive effects of any kind in the long-term (again on the same 7-point scale). Fourth, they indicated which charity they felt more emotionally attached to on the same 7-point scale. Participants in the information condition (but not the other participants) also indicated the extent to which they trust the charity experts' research in this particular case (on a 7-point scale with anchors at *Don't trust at all (1)* and *Trust completely (7)*).

At the end of the study, participants indicated the extent to which they believe that it is possible to measure and compare the effectiveness of different charities (on a 7-point scale with anchors at *Completely impossible (1)* and *Completely possible (7)*), to what extent they believe it is possible to predict the future effectiveness of charities based on their past effectiveness (on the same 7-point scale), and to what extent they want their donations to achieve the greatest amount of good for the largest number of people possible (on a 7-

point Likert scale with anchors at *Strongly disagree (1)* and *Strongly agree (7)*). Finally, they indicated how much they donated to charity during the last year and responded to demographic questions.

3.2 Results and discussion

A series of independent t-tests revealed that the information manipulation increased effective giving in all six tasks (Table 3). Participants generally tended to donate to the less effective option in the no-information condition and to the more effective one in the information condition. An exception was the identifiable victim task, where participants on average donated to the more effective option already in the no-information condition. Another exception was the local (vs. foreign) charity task, where participants on average continued to donate to the ineffective option in the information condition. However, the manipulation still made a significant difference in both of those tasks.

The information manipulation also changed participants' beliefs in all tasks except the identifiable victim task, both with respect to which option saves more lives and with respect to which option has a greater positive impact of any kind in the long-term (Tables 4 and 5). In the information condition, but not in the no-information condition, participants in all six tasks on average correctly identified the charity expected to save more lives. Beliefs about which option saves more lives and beliefs about which option has a greater positive impact of any kind in the long term correlated strongly: $r = .66$ for disaster relief, $r = .63$ for identifiable victim, $r = .88$ for risk aversion, $r = .78$ for overhead costs, $r = .85$ for splitting, and $r = .69$ for local.

The information manipulation also increased the emotional appeal of the more effective option in the disaster relief, risk aversion, overhead, and splitting tasks (Table 6). However, it failed to noticeably boost the emotional appeal of the effective option in the identifiable victim and the local tasks.

Across all six tasks and both the information and the no-information conditions, Participants had a stronger tendency

TABLE 4: Beliefs about which option saves more lives

	No Information	Information	t	df	Cohen's d
Overhead	4.21 (2.26)	4.99 (1.88)	2.66**	195.95	.37
Local	4.11 (1.43)	5.46 (1.68)	6.05***	183.86	.87
Disaster relief	4.19 (1.60)	5.90 (1.30)	8.33***	195.1	1.17
Risk aversion	3.50 (2.38)	5.02 (1.95)	4.97***	195.4	.70
Splitting	4.06 (2.55)	5.49 (1.79)	4.62***	187.07	.64
Identifiable victim	6.10 (1.38)	6.43 (1.02)	1.93	190.29	.27

Note. * $p < .05$, ** $p < .01$, *** $p < .001$.

TABLE 5: Beliefs about which option has greater overall positive effects of any kind.

	No Information	Information	t	df	Cohen's d
Overhead	3.74 (2.10)	4.66 (1.92)	3.22*	196.92	.45
Local	3.89 (1.39)	4.69 (1.53)	3.87***	188.84	.55
Disaster relief	4.56 (1.56)	5.53 (1.39)	4.65***	197.0	.66
Risk aversion	3.57 (2.30)	4.62 (2.05)	3.39**	197.0	.48
Splitting	3.76 (2.36)	5.13 (1.82)	4.61***	192.85	.65
Identifiable victim	5.67 (1.65)	5.97 (1.26)	1.46	192.19	.21

Note. * $p < .05$, ** $p < .01$, *** $p < .001$.

TABLE 6: Emotionally more appealing option.

	No Information	Information	t	df	Cohen's d
Overhead	2.78 (1.70)	3.76 (1.62)	4.13***	196.21	.58
Local	2.38 (1.61)	2.87 (1.89)	1.97	183.84	.28
Disaster relief	3.42 (1.59)	4.04 (1.70)	2.67**	190.91	.38
Risk aversion	3.23 (1.76)	3.99 (1.62)	3.18**	196.84	.45
Splitting	3.57 (1.84)	4.37 (1.59)	3.29**	169.7	.46
Identifiable victim	3.30 (1.75)	3.44 (1.82)	0.56	192.85	.08

Note. * $p < .05$, ** $p < .01$, *** $p < .001$.

to believe that the more effective charity Would save more lives (i.e., had beliefs in line with expert judgements) than to say that they personally would donate to the more effective charity. This was the case in the overhead task, $t(198) = 6.88$, $p < .001$, $d = .49$, the local task, $t(198) = 11.22$, $p < .001$, $d = .80$, the disaster task, $t(198) = 7.43$, $p < .001$, $d = .53$, the risk task, $t(198) = 5.04$, $p < .001$, $d = .36$, the splitting task, $t(198) = 5.07$, $p < .001$, $d = .36$, and the identifiable victim task, $t(198) = 10.86$, $p < .001$, $d = .77$. This confirms the finding from Study 1a that people have preferences for ineffective options in addition to the false beliefs they hold.

Next, we conducted correlation analyses between participants' donation choices in each task and their beliefs about which option would save more lives, their beliefs about the

long-term positive effects of any kind, which option they considered more emotionally appealing, their beliefs about measurability and comparability of the effectiveness of charities, their beliefs about the extent to which future effectiveness can be predicted, and their personal desire to do the most good when donating (Table 7 for the information condition; Table 8 for the no-information condition). The analyses revealed that beliefs about which option would save more lives, beliefs about the long-term positive impact of any kind, and ratings of emotional appeal, were strongly associated with donation choices. Note that belief about saving more lives and belief about greater positive impact of any kind in the long-term were strongly correlated. The desire to do the most good tended to correlate positively with choosing to

TABLE 7: Correlations with donation choice in the information condition.

	Disaster relief	Identifiable victim	Risk	Overhead	Splitting	Local
Belief about saving lives	.56***	.33***	.79***	.72***	.73***	.34***
Long-term impact	.60***	.65***	.84***	.82***	.83***	.55***
Emotionally appealing	.58***	.52***	.63***	.78***	.62***	.74***
Measure & compare	.16	.03	.23*	.29**	.08	.16
Predict effectiveness	.16	.12	.21*	.31***	.13	.19
Do the most good	.41***	.18	.00	.35***	.29***	.33***

Note. * $p < .05$, ** $p < .01$, *** $p < .001$.

TABLE 8: Correlations with donation choice in the no-information condition. (All correlations in first three rows are *** $p < .001$; no others are significant.)

	Disaster relief	Identifiable victim	Risk	Overhead costs	Splitting	Local
Belief about saving lives	.53	.46	.80	.73	.74	.46
Long-term impact	.26	.58	.77	.83	.81	.57
Emotionally appealing	.76	.60	.76	.83	.72	.87
Measure & compare	-.01	-.02	-.08	.04	.07	-.09
Predict effectiveness	-.02	-.18	-.02	-.08	-.01	-.07
Do the most good	-.06	.00	-.15	-.01	-.10	-.03

donate to the effective charity in the information condition but not in the no-information condition. Beliefs about the extent to which the effectiveness of charities can be measured and compared, and beliefs about the extent to which future effectiveness can be predicted from past effectiveness, were not consistent predictors of having donated effectively.

In the information condition, we found that participants on average trusted the charity experts' research. The more they trusted the experts' research, the more they tended to choose the more effective donation option. That was especially true of the disaster relief task ($M = 5.44$, $SD = 1.20$; $r = .42$, $p < .001$), followed by the identifiable victim ($M = 5.43$, $SD = 1.20$; $r = .38$, $p < .001$), the risk aversion ($M = 4.74$, $SD = 1.57$; $r = .60$, $p < .001$), The overhead ($M = 4.77$, $SD = 1.47$; $r = .34$, $p < .001$), the splitting ($M = 5.03$, $SD = 1.51$; $r = .36$, $p < .001$), and the local (vs. foreign) charity ($M = 5.07$, $SD = 1.31$; $r = .17$, $p = .11$) tasks.

Just like in Study 1a, participants on average had a strong desire to do the most good with their donations ($M = 5.87$, $SD = 1.09$). They also tended to believe that it is possible to measure and compare the effectiveness of different charities ($M = 5.51$, $SD = 1.11$) and that it is possible to predict future effectiveness of charities based on their past effectiveness ($M = 5.17$, $SD = 1.13$). Participants' real-world donations in dollars over the past year (i.e., outside of our study) ($M = 498.14$, $SD = 2921.63$, Median = 95) were not associated with their donation choices in our tasks.

The results of this study support our hypothesis that people have multiple misconceptions about charities and that debunking these misconceptions by providing participants with information can make giving more effective. Furthermore, the results confirm our hypothesis that in addition to having misconceptions, people also have preferences for ineffective options.

4 Study 2a: Free choice of charity

In Study 1 we looked at six specific obstacles to effective giving. In Study 2a we explored whether people would choose to donate to the most effective charities in the world when presented with the option of donating to any charity, and if not, why (case 7 in the introduction). Our hypothesis was that people often have false beliefs about which charities are among the most cost-effective, and that even if they knew, they would not necessarily want to donate to them. Thus, we again hypothesized that false beliefs about what charities are most effective and preferences for ineffective options (for specific ineffective charities) together explain why people donate ineffectively.

4.1 Methods

Participants. We recruited 152 US participants from MTurk. Five participants were excluded because they did

not enter the correct URL to a charity website. The final sample consisted of 147 participants (78 females, age $M = 36.21$, $SD = 10.16$).

Procedure. The study had the same 2 (information: yes vs. no) \times 2 (question type: personal-donation vs. effectiveness-belief) between-participants study design as Study 1 had. In the personal-donation conditions, participants were asked: "Suppose we gave you \$1,000 to donate to any charity. You are free to choose any charity out of all real charities in the world. Which charity would you donate to?" In the effectiveness-belief conditions, participants were asked: "Of all the real charities in the world, which one do you believe is the most effective one (i.e., does the most good for the greatest number of people)? In other words, which charity would do the most good if we donated \$1,000 to it?" In the information conditions participants were informed that cost-effectiveness research has shown that only a few charities have been shown to be very effective, and that most well-known charities are much less effective than the most effective charities. They were informed about the evidence-based charity evaluator GiveWell, including a link to the GiveWell website and their ranking of the most effective charities. In all conditions, participants were told that they should take some time to answer the question, and that they should consider conducting some online research into charities.

We contacted five experts to rate the charities that participants chose (besides those in Study 2a also those that participants in the final task in Study 3 chose). The experts were all professional researchers on charity cost-effectiveness and involved in the effective altruism movement. Since most charities' levels of cost-effectiveness have not been carefully researched, we thought that expert judgments were the best available resource. The experts indicated on a 4-point scale from *Clearly implausible* (1) to *Clearly plausible* (4) the extent to which each charity would be a plausible option for a donor who wants to maximize their impact. We then classified charities that got an average rating above 2.5 as *highly effective*, and those below that as *less effective*. In total, 44 out of 107 unique charities listed by participants were classified as highly effective. Agreement among experts was generally high: out of 107 charities, the standard deviation of the ratings was 0 for 37 charities, below 0.5 for 60 charities, and below 1 for 97 charities.

4.2 Results and discussion

The results showed that 0% of participants (0 of 26) in the no-information condition who were asked what charity they would personally donate to chose to allocate the fictitious \$1,000 to a charity that experts judged to be highly effective. Similarly, 0% (0 of 29) of participants in the no-information condition who were asked what charity they believed to be the most effective chose a charity that experts judged to be

highly effective. By contrast, in the information conditions, 41% of participants (21 of 51) who were asked what charity they would personally donate to chose to donate to a charity that experts judged to be highly effective, and 51% (21 of 41) of participants who were asked what charity they believed to be the most effective chose a charity that experts judged to be highly effective. This difference between the two information conditions, however, was not statistically significant, $\chi^2(1) = .56$, $p = 0.45$.

Overall, these results demonstrate that at least in part, people donate ineffectively because they do not know which charities are most effective. Once informed about which charities are the most effective according to research, many say that they would donate to these charities. However, it is noteworthy that a majority of participants who have been informed about which charities experts judge to be highly effective still chose to donate to other charities. This is due, in large part, to the fact that about half of participants did not believe the effectiveness information they were given. Note that this number was higher than in Study 1, where, across the different tasks, participants generally tended to believe the effectiveness information that they were given. One interpretation of these findings is that people tend to reject information which suggests that their favorite charity is not among the most effective.

5 Study 2b

The purpose of Study 2b was to replicate the findings of Study 2a with a simpler study design and a larger sample size. Whereas Study 2a included four conditions, Study 2b only featured two conditions: one where participants were informed about the most effective charities (information) and one where they were not (no-information). Unlike in Study 2a, each participant was asked both which charity they would personally donate to and which charity they believe is most effective. Similarly to Study 1b, this within-participants design allowed us to study why individuals who were given information about the most effective charities still chose to donate to a less effective charity.

5.1 Methods

Participants. We recruited 304 US participants from MTurk. Seven participants were excluded because they did not enter a real charity. The final sample consisted of 297 participants (130 females, age $M = 37.55$, $SD = 10.88$).

Procedure. The study was based on Study 2a with the key difference being that each participant was asked both which charity they would want to donate to and which charity they believed to be the most effective in the world. Accordingly, this study had only two conditions: information and no-information. The two questions participants were asked and

TABLE 9: Breakdown of participant groups, displaying means and standard deviations for perceived relative effectiveness and relative emotional appeal of the two charities.

Condition	Charity they donated to	Charity they believe is most effective	N	Perceived relative effectiveness	Relative emotional appeal
No-information	Ineffective	Effective	5	2.60 (1.14)	6.80 (0.45)
No-information	Ineffective	Ineffective	136	3.90 (1.49)	5.21 (1.42)
Information	Effective	Effective	48	4.04 (1.38)	4.81 (1.39)
Information	Effective	Ineffective	15	4.53 (1.88)	2.47 (1.19)
Information	Ineffective	Effective	12	2.42 (1.08)	6.83 (0.39)
Information	Ineffective	Ineffective	81	3.93 (1.38)	5.22 (1.39)

the information manipulation were identical to those in Study 2a.

After answering the two main questions, participants rated the relative effectiveness of the two charities on a 7-point scale from *Definitely [name of the charity the participant believe is the most effective charity in the world]* (1), *They are equally effective* (4), to *Definitely [name of the charity the participant said they would donate to]* (7). Next, participants were asked to indicate which of these two charities they feel more emotionally attached to, using a corresponding 7-point response scale as in the previous question, *mutatis mutandis*. Finally, participants responded to demographic questions.

5.2 Results and discussion

The results showed that 0% of participants (0 of 141) in the no-information condition and 40% of participants (63 of 156) in the information condition donated a charity that experts judged to be highly effective, $\chi^2(1) = 69.88, p < .001$. And when asked what charity they believed to be the most effective, 4% of participants (5 of 141) in the no-information condition and 38% of participants (60 of 156) in the information condition chose a charity that experts judged to be highly effective, $\chi^2(1) = 50.79, p < .001$. A McNemar’s chi-squared test showed that the proportions of participants who chose an effective charity when asked what charity is the most effective and when asked what charity they would donate to, respectively, did not differ from each other, $\chi^2(1) = 0.03, p = .86$.

A closer breakdown of the responses in the information condition showed the following: 31% of participants chose an effective charity both when asked what charity is the most effective and when asked what charity they would donate to; 8% chose an effective charity when asked the former question, but an ineffective charity when asked the latter question; 10% chose an ineffective charity when asked the former question and an effective charity when asked the latter question; 52% chose an ineffective charity in both cases.

As stated, participants were also asked which charity they

found more emotionally appealing: the one they said they would donate to, or the one they thought is more effective? A one-sample t-test against the mid-point revealed that on average participants felt more emotionally connected to the charity that they said they would donate to compared with the charity that they believed to be the most effective ($M = 5.10, SD = 1.52$), $t(296) = 12.46, p < .001$. However, when asked directly, participants on average considered the charity they would personally donate to to be just as effective as the charity they said is the most effective ($M = 3.88, SD = 1.48$), $t(296) = -1.45, p = .15$. A partial explanation of this finding is that 37% of participants chose the same charity when asked which charity is the most effective and when asked which charity they would donate to. When we excluded those participants, we found that the remaining participants on average tended to think that the charity they believed is the most effective is more effective than the charity that they would personally donate to ($M = 3.46, SD = 1.44$), $t(187) = -5.12, p < .001$. Similarly, the participants who chose two different charities felt more emotionally connected to the charity they donated to than to the charity they believe was most effective ($M = 5.42, SD = 1.58$), $t(187) = 12.33, p < .001$.

Participants who donated to an ineffective charity (as rated by experts) felt more emotionally connected to that charity (relative to the one that they believed to be the most effective) ($M = 5.33, SD = 1.42$) compared with participants who donated to an effective charity (as rated by experts) ($M = 4.25, SD = 1.61$). Table 9 shows a more detailed breakdown of participants based on condition, on whether they donated to an effective charity, and on whether they knew of a highly effective charity. The table reveals, for example, that participants, in both conditions, who donated to an ineffective charity despite knowing that another charity would be more effective felt strongly emotionally attached to the charity they donated to (relative to the charity they knew is more effective).

In sum, in this study we again found that providing participants with information about which charities are the most

effective can increase their interest in donating to them. Just like in Study 2a, we did not find evidence that providing information affects participants' beliefs about which charity is the most effective more than their donation choices in this task. However, notably we did find that participants who donated to ineffective charities were more emotionally attached to them than participants who donated to effective charities. This finding provides some evidence that preferences for emotionally appealing charities may lead people to give to ineffective charities.

6 Study 3: Effective altruists

In Study 3, we studied a sample of effective altruists. The purpose of this study was to test whether a community of people who are dedicated to doing the most good using reason and evidence would know how to donate effectively, and want to do it across all our tasks. While we should expect that effective altruists prefer to give effectively and know how to, this has not yet been demonstrated in a highly controlled setting. It is possible that effective altruists harbor their own personal preferences for ineffective options, or that they have misconceptions about what the most effective options are.

6.1 Methods

Participants. We recruited 230 effective altruists via the Effective Altruism Newsletter and the Effective Altruist Facebook group, of which 193 completed the study. 10 participants were excluded because they stated that they did not identify as effective altruists. 2 participants were excluded because they failed an attention check. The final sample consisted of 181 participants, which we deemed appropriate for the analysis.

Procedure. The study contained the six tasks from Study 1a (presented first) and the task from Study 2a (presented second). As in Study 1a and 2a, this study had a 2 (information: yes vs. no) \times 2 (question type: personal-donation vs. effectiveness-belief) between-subjects design. Both tasks were identical to the ones presented to laypeople in Studies 1a and 2a. Depending on the condition, participants were again asked to indicate either where they would personally donate or which option is the most effective. Again, some participants received additional information about which option is the most effective.

6.2 Results and discussion

We ran the same analyses as in Study 1a and 2a. Our results show striking differences between laypeople and effective altruists. Effective altruists overwhelmingly chose the effective donation option in all four conditions of each task featured in Study 1a. Across those tasks, the two manipulations

made at most a minor difference, and often no difference at all (see Figure 1 and Tables 10–11). This suggests that effective altruists know how to give effectively, and appear ready to act on those beliefs. Effective altruists also stated, to an even greater extent than laypeople, that when they donate to charity, they want to achieve the greatest amount of good for the largest number of people ($M = 6.64$; $SD = 0.64$), $t(584) = 7.19$, $p < .001$.

Next, we looked at how effective altruists navigated the task from Study 2a. In the no-information condition, 85% of participants (34 of 40) chose to donate to a charity that experts judged to be highly effective, and when asked to name a highly effective charity, 100% of participants (55 of 55) produced one that the experts agreed is highly effective, $\chi^2(1) = 6.43$, $p = .01$. In the information condition, 96% of participants (47 of 49) chose to donate to a charity that experts judged to be highly effective, and when asked to name a highly effective charity, 90% of participants (27 of 30) produced one that the experts agreed is highly effective, $\chi^2(1) = 0.33$, $p = .57$.

In sum, we found that effective altruists have a strong preference to give effectively and are also well informed about charities. This was the case across all seven tasks.

7 General discussion

Across five studies, we found that people give to ineffective charities because they lack knowledge about effectiveness and because of preferences for ineffective options. These results support our hypothesis that there is truth to both the belief-based and the preference-based explanation of ineffective giving.

Study 1 revealed five common misconceptions about charity effectiveness: (1) that charities must have low overhead costs in order to be effective; (2) that charities helping individuals in the developing world are not more effective than their counterparts helping individuals in the developed world; (3) that charities focused on recurring health issues are not more effective than disaster relief charities; (4) that it is not more effective to donate to a charity that has a small chance of achieving tremendous success, than to a charity that will have a much lower impact for certain; (5) that splitting donations across multiple charities for the sake of diversification is a generally effective practice. In line with our hypothesis, we did not find any misconception about effectiveness pertaining to the identifiable victim task. Study 2, in turn, showed that people do not know which the most effective charities in the world are.

Across the tasks we tested in Study 1, we found that these misconceptions could substantially be debunked by providing people with information, and that that affected participants' donation choices. These simple information inter-

TABLE 10: Means and standard deviations for each task and condition (Study 3).

Task	Personal-donation		Effectiveness-belief	
	Information	No-information	Information	No-information
Overhead	5.44 (1.91)	6.75 (0.52)	5.72 (1.64)	6.28 (1.35)
Local	6.12 (1.27)	6.51 (1.17)	6.59 (0.99)	6.53 (1.14)
Disaster relief	5.90 (1.26)	6.57 (1.04)	6.07 (0.93)	6.56 (0.98)
Risk aversion	5.71 (1.93)	6.25 (1.04)	6.48 (1.17)	6.91 (0.3)
Splitting	6.15 (1.87)	6.45 (1.19)	6.64 (1.17)	6.47 (1.29)
Identifiable victim	6.46 (1.05)	6.84 (0.42)	6.83 (0.42)	6.66 (1.1)

TABLE 11: ANOVAs with donation choice (from definitely effective to definitely not effective) as the dependent variable and question type (personal-donation vs. effectiveness-belief) and information as independent variables (Study 3).

Task	Question type		Information		Interaction	
	F	η_p^2	F	η_p^2	F	η_p^2
Overhead	1.27	< .001	18.82**	.096	2.92	.016
Local	1.87	.013	1.06	.006	1.65	.009
Disaster relief	0.03	< .001	13.31**	.070	0.29	.002
Risk aversion	11.07**	.074	6.50*	.035	0.11	.001
Splitting	1.60	.013	0.13	.001	1.25	.007
Identifiable victim	0.69	.006	1.06	.006	5.82*	.001

Note. The degrees of freedom for all scenarios were (1, 178). * $p < .05$, ** $p < .001$.

ventions increased the proportion of participants interested in giving to an effective charity. Even in the identifiable victim task, where people did not have any misconception, providing them with information about which option is most effective made them donate more effectively. Likewise, in Study 2, we found strong effects of informing participants about GiveWell’s top-recommended charities: e.g., in Study 2a, it resulted in 41% of participants expressing an interest in donating to one of them.

However, providing information does not eliminate ineffective giving. We found that people often continue to choose the less effective donation option even when they know that they could do more good by choosing a different option. In other words, people have preferences that are at odds with effectiveness. In Study 1, we found preferences for low overhead compared to high overhead charities, for local over foreign charities, for disaster relief charities over charities that address recurring health problems, for charities that have a guaranteed impact over those whose impact is less certain, for splitting donations across charities, and for charities that help identifiable victims over unidentifiable ones. In Study 2b, we found that people who donate to ineffective charities are more emotionally attached to them

than people who donate to effective charities. Some of these preferences have been studied in previous research (Baron & Szymanska, 2011; Caviola et al., 2014; Gneezy et al., 2014), whereas others have not previously been demonstrated. Most notably Berman et al. (2018) have shown that a key reason why people donate ineffectively is that they have subjective preferences for certain charities, and choose to donate to them.

An issue that deserves comment is the relationship between people’s preference for giving effectively, and their preference for giving at all. In all our studies, we presented participants with the choice of two giving options: one effective option, and one ineffective option. Not giving at all was not an option: it was already decided that they donate; the question was where to. But in the real world, people’s decision where to donate is typically not independent of whether to donate. The relationship between people’s motivation to give at all, and their motivation to give effectively, deserves further research (Baron & Szymanska, 2011; Small et al., 2007).

It is also important to distinguish between proximate and ultimate explanations. In this project, we focus solely on the proximate psychological causes of ineffective giving: what

beliefs and preferences people have that can explain why they give to ineffective charities. We did not attempt to give an *ultimate* explanation of why people have these psychological obstacles in the first place. That is a topic for future research (though see some related comments in the next section).

7.1 Misconceptions and preferences are linked

We found that misconceptions about charities and preferences for ineffective options were closely linked. Each misconception had a corresponding preference at odds with effectiveness. The converse was also largely true. (The only exception was that people's preference for giving to identifiable victims over larger numbers of unidentifiable victims was not accompanied by a corresponding misconception.)

The close link between misconceptions and preferences raises the question of causality. Do people hold preferences for ineffective options because of the corresponding misconceptions? Or do people have misconceptions because of their corresponding preferences for ineffective options? Further research is needed to answer this question. We think it is plausible that causality could go both ways. For example, motivated reasoning could lead people to believe that their subjectively preferred charities are unusually effective. If right, that entails that people want to believe that their favorite charities are effective. That may be why we found that the number of people who believed that an ineffective charity is effective was roughly the same as the number of people who wanted to donate to that charity in Study 2. It also fits with our finding that people explicitly stated that they want their donations to achieve the greatest amount of good for the largest number of people. And conversely, the fact that people falsely believe that certain features (such as low overhead costs) predict high cost-effectiveness may cause or strengthen a preference for such features. More generally, people's lack of awareness of the massive disparities in effectiveness separating top charities from average ones may partially explain why they never developed stronger preferences for effectiveness (Caviola et al., 2020).

In Study 3, we found that effective altruists give effectively across all tasks. They had neither misconceptions nor preferences for ineffective options. This finding gives further support to the notion that misconceptions and preferences are closely linked. It is not clear whether their knowledge (their lack of misconceptions) has caused them to have a strong preference for donating effectively, or whether it is rather that their preferences for effectiveness have made them more knowledgeable. Again, it is plausible that the causality goes both ways. Effective altruists may have developed a strong preference to give effectively through learning about the huge differences in effectiveness between charities. At the same time, they may have learned more about charities because of

their pre-existing preference to help effectively.

7.2 Can people become effective donors?

Our research suggests several strategies for increasing the effectiveness of giving in the real world, which future research could explore further.

One approach would be to run campaigns to correct misconceptions about charities. In our studies, such debunkings had substantial effects. However, it is unlikely that the effects would be as strong in the real world, for a number of reasons. First, people might not always believe or understand the information they are given. In particular, in the real world, people may fail to comprehend, remember, and transfer the information they are given to new concrete giving tasks. Second, in more realistic settings, multiple misconceptions usually apply simultaneously. For instance, one may face a choice between multiple charities, the most effective of which helps unidentifiable victims who suffer from recurring health problems in foreign countries. In order to choose such a charity over its competitors, one would need to overcome psychological barriers to effectiveness associated with all of those tasks at once. Therefore, correcting just one or two of them may only make a small difference to the effectiveness of people's donations.

The other approach — to the extent that it is considered ethically desirable — is to change people's preferences. Our research suggests that when asked in the abstract, people want their donations to do the most good for the largest possible number of people. However, when confronted with more concrete donation tasks, they often chose the ineffective option, because of a preference that conflicts with their preference for effectiveness. Can those preferences be changed, so that people donate more effectively? Ultimately, that may depend on whether a social norm emerges to the effect that one ought to donate effectively. At present, there does not seem to be such a norm (Berman et al., 2018); instead the norm is that it is permissible or even mandatory to donate ineffectively, if that is what your heart tells you. Future research could study whether it is possible to change such norms.

The existence of effective altruists demonstrates that it is possible for people to become effective donors. Moreover, their existence suggests a strategy for increasing the effectiveness of giving: study the factors that inspire effective altruists. A deeper knowledge of what motivates effective altruists could conceivably be used to nudge larger segments of the population toward effectiveness.

7.3 Limits to measurability

In our studies, we looked only at cases where cost-effectiveness is relatively easy to measure and compare.

However, in the real world, it is often difficult to measure the effectiveness of charities, to predict their long-term impact, and compare different options. In fact, critics of the effective altruism movement have argued that effective altruists tend to underestimate these difficulties, or else deprioritize effective interventions which are more difficult to measure, such as system-changing interventions (Gabriel, 2017). We will not enter into a detailed discussion of those complex issues here; however, we do think it could be worthwhile for future research to look at cases where cost-effectiveness is harder to evaluate. Will people be more or less willing to donate to supposedly effective charities that focus on systemic change, compared with charities that focus on relatively measurable health interventions? We have not studied such nuances, but they must be confronted to understand effective giving in the real world.

7.4 Conclusion

To sum up, our studies suggest that both lack of knowledge and preferences for ineffective options cause ineffective giving. In the real world, multiple misconceptions, ignorance of what the most effective charities are, and preferences at odds with effectiveness may stand in the way of donors choosing the most effective donation option. These findings may explain why there are so few effective donors.

8 References

- Bagwell, S., de Las Casas, L., van Poortvliet, M., & Abercrombie, R. (2013). *Money for good UK: Understanding donor motivation and behaviour*. London: New Philanthropy Capital.
- Baron, J., & Szymanska, E. (2010). Heuristics and biases in charity. In D. Oppenheimer & C. Olivola (Eds). *The science of giving: Experimental approaches to the study of charity*, pp. 215–236. New York: Taylor and Francis.
- Berman, J., Barasch, A., Levine, E., Small, D. (2018). Impediments to effective altruism: The role of subjective preferences in charitable giving. *Psychological Science*, 29(5), 834–844.
- Brock, J. M., Lange, A., & Ozbay, E. Y. (2013). Dictating the risk: Experimental evidence on giving in risky environments. *American Economic Review*, 103, 415–437.
- Caviola, L., Faulmüller, N., Everett, J. A. C., Savulescu, J., & Kahane, G. (2014). The evaluability bias in charitable giving: Saving administration costs or saving lives? *Judgment and Decision Making*, 9(4), 303–316.
- Caviola, L., Schubert, S., Teperman, E., Moss, D., Greenberg, S., & Faber, N. (2020). People vastly underestimate differences in charities' effectiveness. Under Review.
- Lomborg, B. (2012). *Global Crises, Global Solutions (2012)*, (Ed.). Cambridge University Press.
- Gabriel, I. (2017). Effective altruism and its critics. *Journal of Applied Philosophy*, 34(4), 457–473.
- GiveWell. (2012). Your dollar goes further overseas. Retrieved May 7, 2019, from GiveWell website: <https://www.givewell.org/giving101/Your-dollar-goes-further-overseas>.
- Gneezy, U., Keenan, E. A., & Gneezy, A. (2014). Avoiding overhead aversion in charity. *Science*, 346(6209), 632–635.
- Haynie, J. (2019). Moving beyond the numbers: Evidence that financial ratios don't predict nonprofit outputs. Retrieved from <https://docplayer.net/amp/102543858-Moving-beyond-the-numbers-evidence-that-financial-ratios-don-t-predict-nonprofit-outputs.html>.
- Jamison, D. T., Breman, J. G., Measham, A. R., Alleyne, G., Claeson, M., Evans, D. B., . . . Musgrove, P. (2006). *Disease control priorities in developing countries, 2nd edition*. The World Bank.
- Karlan, D., & Wood, D. H. (2017). The effect of effectiveness: Donor response to aid effectiveness in a direct mail fundraising experiment. *Journal of Behavioral and Experimental Economics*, 66, 1–8.
- Knowles, S., & Sullivan, T. (2017). Does charity begin at home or overseas? *Nonprofit and Voluntary Sector Quarterly*, 46(5), 944–962.
- MacAskill, W. (2015). *Doing good better: Effective altruism and a radical new way to make a difference*. Random House.
- Metzger, L., & Günther, I. (2019). Making an impact? The relevance of information on aid effectiveness for charitable giving. A laboratory experiment. *Journal of Development Economics*, 136, 18–33.
- Nagel, J., & Waldmann, M. R. (2013). Spatial distance, availability of means, and moral obligation judgments. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 35, 3127–3132.
- Sharps, D. L., & Schroeder, J. (2019). The preference for distributed helping. *Journal of Personality and Social Psychology*, 117(5), 954–977.
- Slovic, P. (2007). "If I look at the mass I will never act": Psychic numbing and genocide. *Judgment and Decision Making*, 2, 79–95.
- Small, D. A., Loewenstein, G., & Slovic, P. (2007). Sympathy and callousness: The impact of deliberative thought on donations to identifiable and statistical victims. *Organizational Behavior and Human Decision Processes*, 102(2), 143–153.
- Tollefson, J. (2015). Can randomized trials eliminate global poverty? *Nature*, 524(7564), 150–153.