

TITLE

Donors vastly underestimate differences in charities' effectiveness

AUTHORS

Caviola, L; Schubert, S; Teperman, E; et al.

JOURNAL

Judgment and Decision Making

DEPOSITED IN ORE

31 July 2020

This version available at

<http://hdl.handle.net/10871/122268>

COPYRIGHT AND REUSE

Open Research Exeter makes this work available in accordance with publisher policies.

A NOTE ON VERSIONS

The version presented here may differ from the published version. If citing, you are advised to consult the published version for pagination, volume/issue and date of publication

Donors vastly underestimate differences in charities' effectiveness

Lucius Caviola*[†] Stefan Schubert[‡] Elliot Teperman David Moss[§] Spencer Greenberg
Nadira S. Faber[¶]

Abstract

Some charities are much more cost-effective than other charities, which means that they can save many more lives with the same amount of money. Yet most donations do not go to the most effective charities. Why is that? We hypothesized that part of the reason is that people underestimate how much more effective the most effective charities are compared with the average charity. Thus, they do not know how much more good they could do if they donated to the most effective charities. We studied this hypothesis using samples of the general population, students, experts, and effective altruists in six studies. We found that lay people estimated that among charities helping the global poor, the most effective charities are 1.5 times more effective than the average charity (Studies 1 and 2). Effective altruists, in contrast, estimated the difference to be factor 30 (Study 3) and experts estimated the factor to be 100 (Study 4). We found that participants donated more to the most effective charity, and less to an average charity, when informed about the large difference in cost-effectiveness (Study 5). In conclusion, misconceptions about the difference in effectiveness between charities is thus likely one reason, among many, why people donate ineffectively.

Keywords: cost-effectiveness, charitable giving, effective altruism, prosocial behavior, helping

1 Introduction

People donate large sums to charity every year. US charitable giving amounted to \$410 Billion in 2017: more than 2% of GDP (Giving USA, 2018). Thus, people make significant sacrifices in order to help others. Yet, their help is often much less impactful than it could be. An increasing number of critics have argued that most people donate ineffectively—that they donate to charities which save fewer lives, or in other ways do less good, than the most (cost-)effective charities (Fiennes, 2017; MacAskill, 2015). This contrasts with self-

interested behavior: e.g., people more frequently choose the most effective option when investing (for their own benefit) than when donating (for others' benefit) (Berman, Barasch, Levine & Small, 2018). There is a puzzle here: the *puzzle of ineffective giving*. If people make sacrifices in order to help others, why do they not help others more effectively (Andreoni, 1990; Bergh & Reinstein, 2020; Berman et al., 2018; Caviola, Schubert, Nemirow, 2020; FeldmanHall, Dalgleish, Evans & Mobbs, 2015; Gneezy, Imas, Brown, Nelson & Norton, 2012; Karlan & Wood, 2017; Metzger & Günther, 2015; Null, 2011; Verkaik, 2016)? In this paper, we study a hypothesis which could partly explain the puzzle of ineffective giving: that people underestimate the difference in effectiveness between charities.

In recent years, many scholars have argued that some forms of helping are substantially more effective than others (MacAskill, 2015; Ord, 2013; Singer, 2009, 2015). In particular, the nascent effective altruism movement, pioneered by academic philosophers, has argued that the differences in charity effectiveness are large, and that that is part of the reason why it is so important to donate to the most effective charities (MacAskill, 2015; Ord, 2013; Pummer, 2016; Singer, 2009, 2015). An oft-used example involves two interventions aimed at mitigating blindness that different charities implement. Trachoma surgeries to prevent blindness have been estimated to cost less than \$50, whereas training a guide dog to help a person who is already blind has been estimated to cost \$50,000 (Colby, 2017; Jamison et al., 2006; Ord, 2013). This means that the money spent on training a single guide dog could have funded trachoma surgeries that would have prevented one thousand people from going blind

We thank Robert Wiblin for inspiring the main idea of this paper, Fabienne Sandkühler for her valuable inputs on the manuscript and analysis, and for helping with the data collection, and Jessie Sun, Pablo Stafforini, Philip Trammell and Lewis Ho and members of the Global Priorities Institute seminar for their helpful comments. We also thank the experts who responded to our experts survey and the respondents of our survey of effective altruists.

Lucius Caviola and Stefan Schubert share first authorship. S.G. came up with the main idea, demonstrated the effect in pilot studies, and commented on a draft of the paper. L.C., S.S. and N.S.F. planned the main studies and wrote the paper. E.T. contributed to study 1, D.M. ran study 3 and contributed to study 4, and both commented on a draft of the paper.

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

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

[†]These authors contributed equally.

[‡]Oxford Uehiro Centre for Practical Ethics, University of Oxford, 16–17 St Ebbs St, Oxford OX1 1PT, UK.

[§]Faculty of Education, Canterbury Christ Church University, North Holmes Road, Canterbury, Kent, CT1 1QU, UK

[¶]College of Life and Environmental Sciences, University of Exeter, Washington Singer Building, Exeter EX4 4QG, United Kingdom.

in the first place. So if the goal is to mitigate problems associated with blindness — including both mitigation of current blindness and prevention of potential blindness — then trachoma surgery is far more effective. It seems plausible that most people do not know about these big differences in effectiveness, and this lack of knowledge could reduce their tendency to donate to the most effective charities.

Just like donors, social scientists interested in charitable giving have for the most part neglected the issue of *quality*, or effectiveness, of giving, instead focusing on how much people give—the *quantity* of giving (Bekkers & Wiepking, 2011; van Leeuwen & Wiepking, 2013). But recent years have seen an increased interest in why people give ineffectively (Baron & Szymanska, 2011; Berman et al., 2018; Caviola, Faulmüller, Everett, Savulescu & Kahane, 2014; Karlan & Wood, 2017; Metzger & Günther, 2015; Null, 2011; Verkaik, 2016). There are two main types of explanations for why people do not donate effectively: *motivational* and *cognitive/epistemic* explanations (Berman et al., 2018, Caviola, Schubert & Nemirow, 2020). Motivational explanations posit that people do not want to donate effectively, or that they just have a weak preference for donating effectively, which can be overridden by other preferences. For instance, people may not want to deliberate too much when donating, because deliberate donors are seen as less reliable cooperation partners (Montealegre, Bush, Moss, Pizarro & Jimenez-Leal, 2020; Rand, 2016). And lack of deliberation leads to ineffective giving. Similarly, it has been suggested that people prefer donating to charities that they have a subjective preference for (e.g., cancer charities) over more effective charities (e.g., charities focusing on neglected tropical diseases) and that they do not want to seek information about charities' effectiveness (Berman et al., 2018).

Cognitive/epistemic explanations, on the other hand, say that people donate ineffectively because of cognitive or epistemic shortcomings: because they do not know how to donate effectively. Most of the existing research on cognitive/epistemic explanations of ineffective giving has focused on general cognitive biases. They include insensitivity to the number of people saved (scope neglect; Dickert, Västfjäll, Kleber & Slovic, 2015) and the tendency to focus on what proportion of people one can save, rather than on the absolute numbers (proportion dominance; Fetherstonhaugh, Slovic, Johnson & Friedrich, 1997). These biases cause people to process information about charities and giving opportunities in ways that can lead to ineffective giving.

We suggest an additional epistemic explanation, which has received less attention: that people have *misconceptions* about charitable giving. This *misconception explanation* suggests that the problem is not that people process available information incorrectly, but rather that they do not even start with correct information. In previous work we have already explored this misconception explanation (Caviola et al. 2020). Here, we study a new hypothesized misconception:

that people underestimate the difference in effectiveness between the most effective charities and charities of average effectiveness. Specifically, we focus on charities that help the world's poorest people. We also hypothesize that this misconception affects giving behavior. If people believe that the difference in effectiveness between charities is low, then they may find it less important to donate to the most effective charities. Therefore, it seems plausible that this misconception could be a cause of ineffective giving.

We study our hypotheses about lay misconceptions about charity effectiveness and their effects on giving behavior across five studies. In Studies 1–4, we test whether lay people hold the hypothesized misconception by comparing lay people's and experts' beliefs about the difference in effectiveness between the most effective charities and charities of average effectiveness. In Studies 5, we study whether informing lay people about the differences in effectiveness between the most effective and an average charity affects their giving behavior.

Reports of all measures, manipulations, data (including exclusions), analysis code, and experimental materials in all studies are available for download at: <https://osf.io/k4zft/>.

2 Study 1

Study 1 tested lay people's beliefs about the difference in effectiveness between the most effective charities and charities of average effectiveness that aim to help the world's poorest people.

2.1 Method

Participants. We recruited 210 US American participants via Amazon MTurk. Forty-two were excluded because they did not complete an attention check question correctly or because they said that a charity of average level of cost-effectiveness is more effective than a charity of the highest level of cost-effectiveness. The final sample consisted of 168 participants (89 females, age $M = 38.85$, $SD = 11.63$).

Procedure. Participants first read an explanation of the concept of cost-effectiveness, then they read about two charities:

Some charities that help the world's poorest people are more cost-effective than others. Cost-effectiveness (in our example here) is measured by the number of lives saved. A more cost-effective charity can save more lives than a less cost-effective charity can save with the same amount of money.

Among all the charities that help the world's poorest people.

- Charity A has the highest level of cost-effectiveness
- Charity B has an average level of cost-effectiveness.

Subsequently, participants answered three questions probing their estimates of differences in cost-effectiveness between Charity A and Charity B:

1. *Tipping point*: “Imagine that Charity A receives \$1000, which it uses to save a certain number of lives. How much money would Charity B need to receive, in order to save the same number of lives as Charity A?”
2. *Explicit comparison*: “How many times more cost-effective do you believe Charity A is in comparison to Charity B?”
3. *Cost per life ratio*: “How much do you think it would cost a charity of highest level cost-effectiveness [charity of average level cost-effectiveness] to prevent one person in a poor country from dying? (In US dollars)”

To calculate the cost per life ratio, we divided the participants' estimate of the cost-effectiveness of a charity with the highest level of cost-effectiveness by that of a charity with an average level of cost-effectiveness.

Participants also responded to a post-experimental questionnaire, which included a question about how much they donated to charity last year. Another question asked participants what proportion of charities they believe have the highest level of cost-effectiveness, defined as at least 90% as effective as the most effective charity, and what proportion have an average level of effectiveness, defined as between 10% more and 10% less effective than the average charity. We also asked participants how knowledgeable about charities they believe they are and whether they have heard of Effective Altruism or GiveWell before. Finally, participants responded to demographic questions.

2.2 Results

The results were convergent (see Table 1): the median ratio was 1.5 for *Tipping point*, 2 for *Explicit comparison* and 1.6 for *Cost per life ratio*. Surprisingly, we also found that the median participant believed that the most cost-effective charity can save a life for \$100, and that the average charity can save a life for \$150. Both figures are substantially lower than estimates from the expert charity evaluator GiveWell, which suggest that the most cost-effective charities can save a life for approximately \$2,000 (GiveWell, 2020). There were no differences in the reported cost-effectiveness ratios between donors (i.e., participants who donated at least something to charity in the previous year; 81.5% of them) and non-donors, $t(136) = 0.99, p = .32$.

We found that participants on average thought that 21.74% of charities had the highest level of cost-effectiveness, and

47.89% had an average level of cost-effectiveness. 16 out of 168 participants (9.5%) reported having heard of Effective Altruism, 133 (79.2%) had not heard of it and 19 (11.3%) were not sure. Seventeen participants (10.1%) reported having heard of GiveWell, 134 (79.8%) had not heard of it and 17 (10.1%) were not sure. There were no noteworthy correlations between the dependent variables and the follow-up or demographic measures.

3 Study 2

In Study 2, we tested whether the findings from Study 1 would replicate with a sample that differed from the first sample in terms of nationality, age and education.

3.1 Method

Participants. We recruited 208 students on the University of Oxford campus. Eight were excluded either because they did not complete the attention check question correctly or because they did not complete the full study, leaving a final sample of 200 people (118 females; age $M = 24.31, SD = 7.68$).

Procedure. Participants read the explanation of the concept of cost-effectiveness from Study 1, before answering the *Tipping Point* question from Study 1. We included just one question because there were no noteworthy differences among the three questions we used in Study 1. Participants then answered questions that were part of another, unrelated, study.

3.2 Results

We found that the median response was a cost-effectiveness ratio of 2 (see Table 1), which is similar to the results from Study 1.

4 Study 3

In Study 3, we studied beliefs about differences in cost-effectiveness between charities among members of the effective altruism community (MacAskill, 2015; Singer, 2015). Effective altruists are committed to doing the most good using reason and evidence (e.g., by donating to the most cost-effective charities). We therefore wanted to find out whether their focus on effectiveness was associated with more knowledge of charity effectiveness than people in general.

4.1 Method

Participants. We distributed a questionnaire through an online survey of effective altruists. The survey was dis-

TABLE 1: Estimates of the cost-effectiveness ratio between the most effective and average charities. Participants in four different studies estimated how much more the most effective charities are compared with an average charity, through three different questions: Tipping point, Explicit comparison, and Cost per life ratio. Note that only the Explicit comparison question probed these estimates directly, however. The Tipping Point question concerned how much money Charity B would need in order to save as many lives as Charity A could save with \$1,000, whereas the Cost per life ratio concerned how much it would cost a charity of highest level of cost-effectiveness [average level of cost-effectiveness] to prevent one person in a poor country from dying. In those cases, the estimated cost-effectiveness ratios were inferred from participants' responses.

	M (SD)	Min	25 th %	Median	75 th %	Max
Study 1 (MTurk sample, $n = 168$):						
Tipping point	1.63 (0.49)	1.00	1.30	1.50	2.00	5.00
Explicit comparison	9.32 (23)	1.00	1.50	2.00	2.85	200
Cost per life ratio	597 (7,715)	1.00	1.50	1.60	2.00	100,000
Study 2 (Oxford students, $n = 200$):						
Tipping point	56 (737)	1.00	1.50	2.00	4.00	10,000
Study 3 (Effective Altruists, $n = 861$):						
Tipping point	10^{26} ($3 \cdot 10^{27}$)	1.00	10	50	100	10^{29}
Study 4 (Experts, $n = 45$):						
Explicit comparison	10,857 (66,997)	1.96	14	100	200	450,000

tributed via a number of channels including the Effective Altruist Newsletter, the Effective Altruist Forum and various effective altruist Facebook groups. In total, 2601 participants participated in the survey. The question of interest for this study was placed in the optional part towards the end of the survey, which was completed by 1139 participants. Of these 278 were excluded for either failing a comprehension check or entering a value that was not a number, leaving us with a final sample of 861.

Procedure. Participants read the explanation of the concept of cost-effectiveness from Study 1, before answering the *Tipping Point* question. They also answered a number of questions relating to other studies before and after the questions relating to the current paper.

4.2 Results

The median participant estimated that the most cost-effective charity was 50 times more cost-effective than an average charity — much higher than the estimates from Studies 1 and 2.

5 Study 4

In Study 4 we asked global poverty experts to estimate the differences in cost-effectiveness between charities helping the global poor. We take this expert estimate as the best available estimate of the true difference in cost-effectiveness between the most effective and average charities.

5.1 Method

Participants. We selected experts in areas relevant to the estimation of global poverty charity effectiveness, in areas such as health economics, international development and charity measurement and evaluation. The experts were identified through searches in published academic literature on global poverty intervention effectiveness and among professional organizations working in charity evaluation. We also let respondents identify other relevant experts in the field; so-called “snowball sampling” (Berg, 2006). We recruited 78 participants, but 33 did not complete the survey or indicated that they were not experts, leaving us with a final sample of 45 participants.

Procedure. The experts were given the *Explicit comparison* question from Study. We informed them that the *Explicit comparison* question was designed for lay people, and that it therefore could be underspecified. The participants were therefore asked to interpret this question to the best of their abilities. They also responded to a post-experimental questionnaire.

5.2 Results and Discussion

We found that their median response was a cost-effectiveness ratio of 100 (see Table 1). This number is much higher than the numbers we found in Studies 1 and 2. Even the 25th percentile expert estimate (14) was much higher than the 75th percentile lay person estimate (≤ 4), which suggests that

experts converged on thinking that lay people's estimates are too low.

6 Study 5

Studies 1–4 together provided evidence that lay people do indeed underestimate how much more effective the most effective charities are compared to the average charity. In Study 5, we studied whether correcting this misconception would lead to more effective giving. Our hypothesis was that when people were correctly informed about the large differences in charities' effectiveness, they would donate more to a highly effective charity, and less to a charity of average cost-effectiveness.

6.1 Method

Participants. We recruited 423 US American participants from MTurk (185 females, age $M = 38.40$, $SD = 12.08$). None were excluded.

Procedure. Participants were divided into three conditions: control, 1.5 ratio and 100 ratio. In the 1.5 and 100 ratio conditions, participants were informed that “in a recent report, researchers concluded that the most cost-effective charities are [1.5/100] times more effective than a charity of average level of cost-effectiveness”. In the control condition, no such information about the cost-effectiveness ratio was given. Next, participants of all conditions were presented with two hypothetical charities (Charity A and Charity B). The charities were described to focus on “helping people in an African town who are at risk of a mosquito-borne infection called Chikungunya [West Nile Virus] by distributing the medicine Doxycycline [Lariam].” The charities' descriptions were randomized, i.e., either Charity A focused on Chikungunya using Doxycycline and Charity B on West Nile Virus using Lariam, or vice versa. Participants were informed that “Among all the charities that help the world's poorest people, Charity A has the highest level of cost-effectiveness and Charity B has an average level of cost-effectiveness.”

Next, participants were asked how they would allocate a hypothetical donation of \$100 between the two charities. They could choose between two options: 1) \$100 to Charity A and \$0 to Charity B, 2) \$60 to Charity A and \$40 to Charity B. We hypothesized that participants would be more likely to give the full amount to the highly effective charity when told the effectiveness ratio was 100, whereas they would still want to give at least some amount to the average charity when told the effectiveness ratio was only 1.5. We hypothesized that people in the control condition would give similar responses as in the 1.5 ratio condition because their implicit assumption

about the effectiveness ratio would roughly match factor 1.5, in line with in our previous studies.

6.2 Results and Discussion

Our hypothesis was supported. The percentage of participants who gave the full amount to the highly effective charity was 37.3% (53 of 142) in the control condition, 38.0% (52 of 137) in the 1.5 ratio condition, and 55.6% (80 of 144) in the 100 ratio condition. A logistic regression revealed that the increase in percentage who gave the full amount to the highly effective charity from the 1.5 ratio condition to the 100 ratio condition was statistically significant ($Exp(B) = -0.71$, $z = -2.94$, $p = .003$), whereas there was no significant difference between the 1.5 ratio condition and the control condition ($Exp(B) = .03$, $z = 0.11$, $p = .91$).

We know from previous research that when presented with multiple charities, people have an inclination to split their donations (Sharps & Schroeder, 2018). They tend to split their donations even in cases where they could do more good by giving everything to the more effective charity (Baron & Szymanska, 2011; Caviola et al., 2020). Our finding that many participants split their donations between a highly effective and an average charity is in line with that. However, as we saw, the tendency to split donations was reduced when the difference in effectiveness between these two charities was said to be large.

One possible explanation is that people's desire to split (e.g., due to fairness concerns) is outweighed by a stronger desire for effectiveness when the difference in effectiveness is large. Another explanation could be that people falsely assume that the most effective allocation is to give X times more (i.e., either 1.5 or 100) to the more effective charity if the charity is X times more effective. Thus, they may fail to realize that the marginal value of every additional dollar is always higher when given to the more effective charity, even in cases where the difference in effectiveness between the two charities is small (Baron & Szymanska, 2011). Irrespective of what explains the effect, this study shows that at least in certain donation tasks, informing people about the large difference between charities' effectiveness can increase effective giving.

More research is required to determine how much the effect of informing people about the large difference in charities' effectiveness generalizes across contexts. It is unclear, for example, if people would also give more effectively if they were considering two charities that focus on completely different causes and are thus more difficult to compare using a single metric. In studies reported in the Supplement, we found mixed evidence for such cases.

7 General Discussion

We found that lay people substantially underestimate the differences in cost-effectiveness between charities, compared with expert estimates (Studies 1–4). Thus, people do indeed hold a misconception about how much greater impact they can have by donating to one of the most effective charities, rather than to an average charity. We also found that it is possible to dispel this misconception: if people are informed that the cost-effectiveness ratio is larger than they thought, they believe it. Furthermore, we found that dispelling this misconception can lead to more effective giving. People are more inclined to give their whole donation to a highly effective charity, and less inclined to split it with a charity of average effectiveness, when informed of the large difference in charities' effectiveness (Study 5).

Our finding adds to the list of misconceptions and preferences for ineffective charities that were revealed in previous research (Baron & Szymanska, 2011; Berman et al., 2018; Caviola et al., 2020). They included, e.g., misconceptions about the effectiveness of disaster charities and local charities, and misconceptions about the effectiveness of splitting donations between charities. They also included corresponding preferences for disaster charities and local charities, and for splitting one's donations, even when participants were informed that donating to those charities or in that way is less effective. In Study 5, most of these preferences and misconceptions did not play any role, with the exception that misconceptions about and preferences for splitting probably made people donate less to the most effective charity than they otherwise would have. But when people are presented with real donation opportunities in the real world, many of these misconceptions and preferences will often stand in the way of effective giving. That could reduce the effect of dispelling the misconception that the difference in effectiveness between charities is low.

On the other hand, there are also considerations pointing in the opposite direction, suggesting that dispelling this particular misconception could have greater effects than our studies suggest. That is because people who truly grasp the significance of the large difference in charities' effectiveness may, besides changing their own giving behavior, also be motivated to influence other donors. They see how much potential impact is lost through ineffective donations, and may therefore try to encourage more effective giving. In fact, members of the effective altruism movement are doing precisely that.

7.1 Why people underestimate the difference in effectiveness between charities

Our studies focused on demonstrating that people erroneously believe that the difference in effectiveness between the most effective and the average charities is small and what

effect this misconception has on their charitable giving. This leaves open the question of *why* they do have this misconception.

One possibility is that people conflate overhead with effectiveness, and (probably correctly) think that the difference between charities' overhead expenditures are relatively modest. Similarly, it is possible that most people do not consider the fact that charities rely on different interventions (e.g., different types of medications) and focus on different problems (e.g., different diseases). Since much of the difference in effectiveness between charities stems from them employing different interventions and addressing different problems, this may lead to them underestimating that difference.

Another possibility is that people make a mistaken analogy between charities and for-profits. The difference in price for two similar products is rarely higher than factor 1.5, because market pressures usually lead to cost-effectiveness optimization (Mankiw, 2011). A seller whose products are 100 times more expensive than other similar products would quickly go bankrupt. Thus, if similar market pressures were operative on the charity "market", one might have expected the differences in effectiveness between charities to be low. However, since donors do not consistently donate more to more cost-effective charities (Karlan & Wood, 2017), there is no equivalent market pressure in the charity sector (i.e., the charity market is not efficient), and as a result there is much less cost-effectiveness optimization. This leads, in turn, to huge differences in cost-effectiveness between charities. Hence, the analogy fails.

Yet another possibility is that the misconception about differences in effectiveness between charities is related to another misconception: people's *overestimation* of how effective charities are in absolute terms, which we observed in Study 1. Since they believe that an average charity is already extremely effective at saving lives, they might believe that it is simply not possible for the most effective charities to be substantially more effective. This leads to a suppressed estimate of the difference in effectiveness between charities. On this view, the two misconceptions are related.

To further understand why people underestimate the difference in effectiveness among charities, it may be useful to study to what extent people have similar misconceptions in other domains, such as regarding the relative effectiveness of different private companies. Such studies could teach us whether the effect that we've found is unique to the charitable domain, or whether it rather reflects more domain-general patterns.

7.2 Limitations and future research

One limitation of our research is that we studied only charities focused on helping the global poor. It is possible that people would have different beliefs about charities working in different cause areas, such as climate change or animal

welfare. Future research could therefore study beliefs about charities working within those cause areas, as well as beliefs about differences in effectiveness between charities working in different cause areas.

Another limitation of our research is that our studies relied only on hypothetical donations. Future research could test whether informing people about the vast difference between charities' effectiveness can increase effective giving in the real world.

7.3 Conclusions

Our research shows that lay people believe that the difference in effectiveness between charities is much smaller than it actually is. We also found that informing them about the median expert estimate of that difference makes them donate more effectively in a donation task. If this finding generalizes, then information about differences in effective giving could be a tool to make charitable giving more effective. Hence more research on the effects of debunking misconceptions about the difference in effectiveness between charities would be valuable.

References

- Andreoni, J. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *The Economic Journal*, *100*(401), 464–477. <https://doi.org/10.2307/2234133>.
- Baron, J., & Szymanska, E. (2011). Heuristics and biases in charity. *The Science of Giving: Experimental Approaches to the Study of Charity*, 215–235.
- Bekkers, R., & Wiepking, P. (2011). A literature review of empirical studies of philanthropy: Eight mechanisms that drive charitable giving. *Nonprofit and Voluntary Sector Quarterly*, *40*(5), 924–973. <https://doi.org/10.1177/0899764010380927>.
- Berg, S. (2006). Snowball sampling — I. In *Encyclopedia of Statistical Sciences*. <https://doi.org/10.1002/0471667196.ess2478.pub2>.
- Bergh, R., & Reinstein, D. (2020). Empathic and numerate giving: The joint effects of victim images and charity evaluations. *Social Psychological and Personality Science*. <https://doi.org/10.1177/1948550619893968>.
- Berman, J. Z., Barasch, A., Levine, E. E., & Small, D. A. (2018). Impediments to effective altruism: The role of subjective preferences in charitable giving. *Psychological Science*, *29*(5), 834–844. <https://doi.org/10.1177/0956797617747648>.
- 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., & Nemirow, J. (2020). The many obstacles to effective giving. *Judgment and Decision Making*, *15*(2), 159–172.
- Colby. (2017). How much does a guide dog cost? *Puppy In Training*. Retrieved from <https://puppyintraining.com/how-much-does-a-guide-dog-cost/>.
- Dickert, S., Västfjäll, D., Kleber, J., & Slovic, P. (2015). Scope insensitivity: The limits of intuitive valuation of human lives in public policy. *Journal of Applied Research in Memory and Cognition*, *4*(3), 248–255. <https://doi.org/10.1016/j.jarmac.2014.09.002>.
- FeldmanHall, O., Dalgleish, T., Evans, D., & Mobbs, D. (2015). Empathic concern drives costly altruism. *NeuroImage*, *105*, 347–356. <https://doi.org/10.1016/j.neuroimage.2014.10.043>.
- Fetherstonhaugh, D., Slovic, P., Johnson, S., & Friedrich, J. (1997). Insensitivity to the value of human life: A study of psychophysical numbing. *Journal of Risk and Uncertainty*, *14*(3), 283–300.
- Fiennes, C. (2017). We need a science of philanthropy. *Nature*, *546*(7657), 187–187. <https://doi.org/10.1038/546187a>.
- GiveWell. (2020). GiveWell's Cost-Effectiveness Analyses. Retrieved June 2, 2020, from <https://www.givewell.org/how-we-work/our-criteria/cost-effectiveness/cost-effectiveness-models>.
- Giving USA. (2018). Giving USA 2018. Retrieved April 4, 2019, from <https://givingusa.org/tag/giving-usa-2018/>.
- Gneezy, A., Imas, A., Brown, A., Nelson, L. D., & Norton, M. I. (2012). Paying to be nice: Consistency and costly prosocial behavior. *Management Science*, *58*(1), 179–187.
- Jamison, D. T., Breman, J. G., Measham, A. R., Alleyne, G., Claeson, M., Evans, D. B., ... & Musgrove, P. (Eds.). (2006). *Disease control priorities in developing countries* (2nd ed.). Retrieved from <http://www.ncbi.nlm.nih.gov/books/NBK11728/>.
- 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. <https://doi.org/10.1016/j.jsocec.2016.05.005>.
- MacAskill, W. (2015). *Doing good better: Effective altruism and a radical new way to make a difference*. Guardian Faber Publishing.
- Mankiw, N. G. (2011). *Principles of economics*. Cengage Learning.
- Metzger, L., & Günther, I. (2015). Making an impact? The relevance of information on aid effectiveness for charitable giving. A laboratory experiment. *ETH Zurich*. <https://doi.org/10.3929/ethz-a-010492382>.
- Monteleone, A., Bush, L., Moss, D., Pizarro, D., & Jimenez-Leal, W. (2020). Does maximizing good make

- people look bad? Manuscript under review. <https://psyarxiv.com/2zbax/>.
- Null, C. (2011). Warm glow, information, and inefficient charitable giving. *Journal of Public Economics*, 95(5), 455–465.
- Ord, T. (2013). The moral imperative toward cost-effectiveness in global health. *Centre for Global Development*.
- Pummer, T. (2016). Whether and where to give. *Philosophy and Public Affairs* 44(1), 77–95.
- Rand, D. G. (2016). Cooperation, fast and slow: Meta-analytic evidence for a theory of social heuristics and self-interested deliberation. *Psychological Science*, 27(9), 1192–1206. <https://doi.org/10.1177/0956797616654455>.
- Sharps, D., & Schroeder, J. (2018). The preference for distributed helping. *Journal of Personality and Social Psychology*, 117(6), 954–977.
- Singer, P. (2009). *The life you can save: Acting now to end world poverty*. Random House Publishing Group.
- Singer, P. (2015). *The most good you can do: How effective altruism is changing ideas about living ethically*. Yale University Press.
- van Leeuwen, M. H. D., & Wiepking, P. (2013). National Campaigns for Charitable Causes: A Literature Review. *Nonprofit and voluntary sector quarterly*, 42(2), 219–240. <https://doi.org/10.1177/0899764012467084>.
- Verkaik, D. (2016). *Do donors really care about impact information? A dual process account*. https://www.researchgate.net/publication/318722670_Do_Donors_Really_Care_About_Impact_Information_A_Dual_Process_Account.