

Review

The Psychology of (In)Effective Altruism

Lucius Caviola,^{1,*} Stefan Schubert,² and Joshua D. Greene^{1,3}

The most effective charities are hundreds of times more effective than typical charities, yet few donors prioritize effectiveness. Why is that? How might we increase the effectiveness of charitable giving? We review the motivational and epistemic causes of (in)effective giving. Many donors view charitable giving as a matter of personal preference, which favors decisions based on emotional appeal rather than effectiveness. In addition, while many donors are motivated to give effectively, they often have misconceptions and cognitive biases that reduce effective giving. Nearly all research on charitable giving interventions focuses on increasing donation amounts. However, to increase societal benefit, donation effectiveness is likely to be more important. This underscores the need for research on strategies to encourage effective giving.

The Importance of Effectiveness

A US\$100 donation can save a person in the developing world from trachoma, a disease that causes blindness [1]. By contrast, it costs US\$50 000 to train a guide dog to help a blind person in the developed world. This large difference in impact per dollar is not unusual. According to expert estimates, the most effective charities are often 100 times more effective than typical charities [2]. In the USA, people donate approximately US\$450 billion [2% of gross domestic product (GDP)] each year [3], but relatively little of this supports the most effective charities – the ones that save the most lives, or improve lives the most, with each dollar (Box 1).

Most research on charitable giving focuses on the amounts that donors give [4]. However, if the societal goal of charitable giving is to improve human (or animal) well-being, it is probably more important to focus on the effectiveness of giving. This is because the variation in charity effectiveness is much larger than the variation in donation amounts that a donor is likely to consider: you can double your impact by doubling the amount that you give to typical charities, but you can multiply your impact by a factor of ten, 100, or even 1000 by choosing to support more effective charities [2]. In recent years, the effective altruism movement has emphasized the enormous variation in charity effectiveness and, consequently, the importance of giving more effectively [5–7] (Effective Altruism).

In addition to its potential social impact, research on the psychology of (in)effective giving is also theoretically significant. It addresses longstanding questions in moral psychology, such as: are humans truly altruistic and if so, to what extent and in what contexts [8]? What motivations and cognitive processes support altruism?

Here, we review recent research on the causes of (in)effective altruism, with a focus on charitable giving. (Other topics include volunteering and career choice.) We present a framework that distinguishes between: (i) motivational and (ii) epistemic influences on effective altruism. We also discuss (iii) interventions that could encourage more effective altruism.

Motivational Obstacles to Effectiveness**Subjectivity of Giving**

Across many domains of economic life, people care about cost-effectiveness. Consumers want the most bang for their buck. Investors want higher returns. Citizens are discouraged by

Highlights

Effective altruism is a philosophy and social movement that advocates using the most effective, evidence-based strategies to benefit others. Here we focus on charitable giving, a domain in which ordinary people can have a large impact.

Most behavioral research on charitable giving focuses on donation amounts, but the impact of giving depends more on the effectiveness of the charities people support than on how much they give. We review recent research on the factors that promote (in)effective giving.

There are motivational and epistemic obstacles to effective giving: People are often drawn to less effective charities, and to the extent that people want to give effectively, they typically do not know how to do it.

We discuss strategies to encourage effective giving. Several strategies are feasible and warrant further research, as the potential social benefits are large.

¹Department of Psychology, Harvard University, Cambridge, MA, USA²Uehiro Centre for Practical Ethics, University of Oxford, Oxford, UK³Center for Brain Science, Harvard University, Cambridge, MA, USA*Correspondence: lcaviola@fas.harvard.edu (L. Caviola).

ineffective public spending. However, when it comes to charitable giving, donors seem to care less about effectiveness. Why is that?

Many people perceive charity as a matter of personal choice, such that donations need not be guided by objective measures of effectiveness [9]. For many, a donor choosing a charity is more like a gourmand choosing a restaurant than a doctor prescribing a treatment. When presented with charities addressing different causes, people tend to choose the charities that have the most personal appeal, even when informed that other options are more effective [9]. Moreover, third parties approve of donation decisions based on emotional connection more than they approve of donation decisions based on objective measures of effectiveness [9]. The idea that charity choice is a matter of subjective preference is widespread, but some philosophers have challenged this idea, arguing that if one is going to help others, one is obliged to do so as effectively as possible (Box 2). While addressing this philosophical controversy is beyond the scope of this review, as a psychological matter the perceived subjectivity of giving may be the most fundamental obstacle to effective giving.

If choosing among charities is perceived as merely an expression of personal preference, donors may see little reason to counteract the misconceptions and cognitive biases that impede effective giving. Likewise, there is little need for them to conduct research or seek expert advice to guide their decisions. In a survey of 3000 American donors [10], only 33% said that they researched the charities they considered and only 9% reported giving based on evidence of relative effectiveness. This contrasts starkly with people's behavior as consumers and investors, where people attend closely to product reviews and seek out financial advice at great expense [11].

People may view donation decisions as subjective because they do not feel responsible or accountable [12] for the effects of their donations. Many people view donating to help strangers as nonobligatory ('supererogatory') [13], which suggests that there is no obligation to give in a particular way. However, there are some contexts in which people feel more obliged to consider effectiveness. For instance, in preliminary research we found that people feel more obliged to maximize effectiveness when they are the only possible helpers and are therefore solely responsible for the outcome [14]. However, individual donors are typically not the only ones who can help, making effectiveness seem less important.

Narrow Affective Motivation

Donors' reliance on subjective preferences raises the question: what produces these preferences? Research indicates that charitable giving, and helping behavior more generally, is typically driven by emotional motivators such as empathy (or sympathy) for victims [15–17] or the positive satisfaction of personally provisioning a good (i.e., 'warm glow' [18]). However, these affective responses are not necessarily triggered by opportunities to help effectively. On the contrary, they often hinder more effective giving [19]. We refer to these as 'narrow' affective motivations, contrasting them with motivations that reflect a broader integration of information about available options and their likely consequences [20].

Personal Connection

People are more inclined to support causes to which they feel personally connected or that they find personally meaningful [9]. For example, donors in wealthy nations are more likely to be personally affected by cancer than malaria, leading to greater support for charities focusing on this more personally relevant disease [21]. Moreover, they are unlikely to change their minds if informed that charity experts consider malaria charities more effective at saving lives. Likewise,

people are often drawn toward charities that help victims suffering from the same diseases or misfortunes as their friends or loved ones [22]. Similarly, people tend to prefer animal charities that help familiar companion animals, such as dogs, over charities that help farm animals, such as pigs, although charities that help farm animals tend to reduce animal suffering more effectively (Animal Charity Evaluatorsⁱⁱ). Finally, people are typically loyal to the charities with which they already have a relationship as a donor [10], a form of *status quo* bias [23].

Narrow Moral Circle

People tend to favor proximate beneficiaries over distant ones. First, there is spatial distance, which typically coincides with social distance: people are more willing to help others the more they feel socially connected with them [24]. Donors in the developed world, for example, prefer local over foreign charities although charities working in distant poor countries tend to be more effective [13,25–27]. Second, there is biological distance: most people prefer to help humans over animals, even when controlling for perceived cognitive capacity and capacity for suffering [28,29]. The tendency to favor humans over animals is a stable trait that correlates with tendencies that are widely regarded as biases (e.g., racism, sexism) and is reflected in donation behavior (helping humans vs animals) [28]. Third, there is temporal distance: people prefer to help current generations over future generations, as reflected in the quip ‘What has posterity ever done for us?’. While strategies to safeguard future generations are by their nature uncertain, they may still be highly effective, as the inhabitants of the future may far outnumber the living. A strategy with a 1% chance of preventing human extinction could save vast numbers of lives, making it a rational bet from an impartial perspective [30]. Although the value of such investments is difficult to quantify, there is evidence that people’s intuitive thinking about extinction events is irrationally myopic, focused on the immediate consequences of catastrophic events rather than their long-term consequences [31].

Scope Neglect

Few would say that the 100th life one can save is worth less than the first life one can save, yet people’s altruistic motivations do not scale proportionately with the number of beneficiaries [32]. Rather, the perceived value of each additional life saved is reduced once large numbers have been saved. This has been called ‘psychophysical numbing’ [33–35]: compassion fades with increasing numbers of victims. In one study, people were willing to donate about the same amount to help either 2000, 20 000, or 200 000 victims [36,37]. This neglect of scope may follow from our tendency to represent values using domain-general mechanisms that evolved to represent the value of goods such as food, which (unlike human lives) exhibits steeply declining marginal value [38].

Scope neglect is accentuated for single identifiable victims, which often draw more support than large numbers of ‘statistical’ victims [39–43]. This is illustrated by the famous case of ‘Baby Jessica’ who, while trapped in a well in Texas, received over US\$700 000 in donations [44]. Furthermore, adding statistical information about victims to profiles of identifiable victims can even reduce people’s willingness to donate [45].

Prioritization Aversion

Donating effectively requires, above all else, the prioritization of charities focused on high-impact causes [5]. The most effective charities tend to work on causes that are relatively: (i) large in scope (more individuals are affected); (ii) neglected (do not receive much support), and (iii) tractable (progress can be made more easily). Charities addressing large, neglected problems for which there are (or may soon be) scalable solutions can have an impact that is orders of magnitude larger than that of typical charities.

Prioritizing some causes over others implies prioritizing some people over others – not because they are inherently less worthy, but as a form of philanthropic ‘triage’. People are averse to making such trade-offs, especially when they involve human lives [46]. Consequently, people consider it unfair to deprioritize certain victims because others can be helped more effectively [47,48]. This is one reason why many donors prefer to split their donations across charities [49], even when some recipient charities are much less effective than others [21,50]. However, this can reduce the overall impact of one’s giving.

While people resist the deprioritization of whole causes based on effectiveness considerations, they do consider it important to avoid charities that they perceive as wasteful. That is, they deprioritize charities that use their available resources in a suboptimal way to serve their cause. For example, people consider it wrong to support a charity that relies on a relatively ineffective medication if another charity addresses the same cause with a more effective intervention [14]. Many donors are especially concerned about supporting organizations with excessively high overhead costs [51,52] (see Overhead Myth later). These tendencies indicate that donors do care about effectiveness. When asked directly, donors often identify effectiveness as a primary motivation for giving [10,21,53]. However, donors’ concern for effectiveness is rather narrow, focused on avoiding wastefulness in the service of a given cause rather than on using effectiveness as a criterion for choosing between causes [9].

Character and Reputational Benefit

One powerful driver of prosocial behavior is reputational benefit [54–56]. Giving to charity can signal good character to potential cooperation partners, but the effectiveness of one’s donations may not influence the strength of that signal, as social rewarders pay little attention to effectiveness [57]. This may be because effectiveness has historically been difficult to track, which puts a reputational premium on prosocial actions that are well-defined and highly observable [57]. This favors visible personal sacrifice over social benefit [58] and donations based on mutually salient emotional factors [59] rather than complex calculations [60]. Consistent with this, people whose donations are based on deliberation rather than empathy are viewed less positively [61]. Under prevailing norms, donors have relatively little reputational incentive to give effectively.

Epistemic Obstacles to Effectiveness

In addition to motivational factors, epistemic factors prevent people from giving effectively, even when they might be motivated to do so.

Overhead Myth

People tend to avoid charities with high overhead (administrative) costs [26,62]. This is partly because they incorrectly believe that overhead ratio tracks cost-effectiveness [62], such that money is always better spent directly on beneficiaries. While charities can overspend on overhead, research indicates that higher overhead costs are not associated with lower cost-effectiveness [62–64]. On the contrary, to maximize effectiveness charities require competent staff, good infrastructure, and self-evaluation, all of which increase overhead costs. Donors’ focus on low overhead creates perverse incentives for charities, leading them to minimize their administrative costs at the expense of effectiveness [63]. Research shows that the ‘overhead myth’ can be partly dispelled. When people are informed that overhead ratio and cost-effectiveness are not the same, many donate to more effective charities, ignoring overhead costs [21,62]. However, the effects of such information may be limited because overhead aversion has a motivational component as well. Spending on overhead can reduce the warm glow of giving [18,52] as the relationship between donor and recipients becomes

more indirect. A more general strategy for charities to mitigate the effects of overhead aversion is to have major philanthropists cover overhead costs, allowing ordinary donors to cover direct costs [52].

Quantifiability Skepticism

Many people doubt that charity effectiveness can be quantified and compared. While people readily accept that one can compare the effectiveness of charities that work on the same problem (e.g., malaria) [21], they have doubts about comparing the effectiveness of charities that work on different problems. There are, however, tools developed by economists that can be used to quantify and compare charity effectiveness. For instance, health economists conduct cost-effectiveness analysis using quality-adjusted life years (QALYs) [65] to quantify the effectiveness of alternative health interventions (Box 1) [66]. While practically challenging, cost-effectiveness analysis allows comparisons of charities addressing different causes (e.g., malaria prevention versus education). More research is needed to investigate people's beliefs about the quantification and comparability of charity effectiveness.

Innumeracy

Maximizing the effectiveness of one's charitable giving may require complex quantitative reasoning, or at least an understanding of certain basic statistical and economic concepts such as probability and expected value. Many people lack such knowledge or fail to transfer it from other domains (e.g., investing) [67]. For example, many donors incorrectly believe that a charity that saves one life with certainty is more effective than a charity that has a 10% chance of saving 100 lives, although the latter can be expected to save ten times more lives over time [21]. Another important but widely misunderstood concept is marginal value, considering the good that an additional donation can do [26]. When one charity is substantially more effective than the alternative charities,

Box 1. How Is Charity Effectiveness Measured?

The effectiveness of charities focusing on health (e.g., treating trachoma; see The Importance of Effectiveness in the main text) can be measured in the same ways that health economists measure the effectiveness of medical treatments [65,66]. Here, researchers calculate the amount of money needed to produce one unit of a given positive outcome (e.g., dollars per life saved). A common outcome measure in health economics is the quality-adjusted life year (QALY) [65,66]. The QALY accounts for the duration of life saved and the health quality of a life (including both psychological and somatic aspects). QALY estimates are rough and imprecise, but that is inevitable for any measure that attempts to quantify important aspects of subjective experience, such as the difference between living with and without a specific painful ailment. By comparing the cost per QALY across interventions, the cost-effectiveness of different charities can be compared.

There is no consensus on how to measure the effectiveness of charities that address other causes, such as the arts, education, or homelessness. Some propose generalizing QALYs by replacing them with well-being-years (WELLBYs) [99,111]. The WELLBY measure aims to incorporate all aspects of well-being, not just health outcomes. While the WELLBY measure reflects an appealingly broad understanding of effectiveness, it does not eliminate the challenges inherent in comparing charities that address very different causes. Nevertheless, cost-effectiveness analysis using WELLBYs or QALYs focuses attention on the number of people impacted and the duration of impact. It provides a general framework within which dimensions of well-being can be weighted in a consistent way. Because differences in charity effectiveness are so large, even highly approximate measures of absolute effectiveness can yield meaningful assessments of relative effectiveness.

Comprehensive measures of impact must account for indirect and long-term effects in addition to short-term, direct effects. For example, curing children in the developing world of parasitic worms may allow them to go to school, with long-term indirect positive effects on education and the local economy [100,101,112]. Tracking indirect and long-term effects poses no special theoretical difficulties, since such effects can be incorporated into cost-effectiveness analysis, but is an additional practical challenge.

In some cases, estimates of effectiveness will depend on ethical judgments, such as the value placed on present versus future generations [102,103], humans versus animals [104], or bringing new people into existence [105]. The nascent field of global priorities research [106] gathers insights from philosophy, psychology, and economics to address these questions.

one maximizes impact by donating exclusively to the most effective charity (assuming there is no substantial decline in marginal value) [68]. Many people incorrectly believe that it is more effective to split their donations across charities that vary in effectiveness [21,26,69].


People are also misled by ‘pseudoinefficacy’: if a vast number of people are threatened, people often view saving a small number of people as just a ‘drop in the bucket’ and therefore not valuable [70]. In other words, donors attend to the proportion of victims they can save rather than the absolute number [33]. This can discourage donors from supporting solutions that, while highly effective, address problems of broad scope.

Underestimation of Effectiveness Variance

Most people underestimate the variation in charity effectiveness. Experts believe that the most effective charities working on global poverty are about 100 times more cost-effective than typical charities in that domain [2]. Laypeople, by contrast, estimate that the most effective charities are only about 1.5 times more effective than average [2]. People assume that the distribution of charities by effectiveness is approximately normal (like the distribution of adult human heights) when it more closely follows a power law distribution, such that a small number of charities are vastly more effective than typical charities, much as the largest urban areas are vastly larger than typical urban areas (Figure 1). One explanation for these vast differences in effectiveness is that the charity market is, from an effectiveness perspective, highly inefficient. Donors do not penalize charities that do less good than others in the way consumers penalize companies that offer higher prices or lower quality products. As a result, many relatively ineffective charities remain in operation without improving.

Ignorance About the Most Effective Charities

Effective altruist researchers such as those at GiveWell^{iv} aim to identify the world’s most effective charities. They have converged on a relatively small number of stand-out charities that focus on supporting the world’s poorest people (GiveWell’s focus), improving conditions at factory farms, and safeguarding humanity’s long-term future (e.g., from global pandemic risks). Most



Trends in Cognitive Sciences

Figure 1. Cost-Effectiveness Distribution of Interventions in Global Health. There are large differences in cost-effectiveness across health interventions. This figure displays the cost-effectiveness estimates for 108 health interventions, arranged from least cost-effective to most cost-effective. Cost-effectiveness is here measured in disability-adjusted life years (DALYs) per US\$1000. Originally reported in [1], based on the Disease Control Priorities in Developing Countries (DCP2) Report [110]; figure adapted by 80 000 Hoursⁱⁱ, used here with permission.

people are unaware of these charities [21], which partly explains why few people support them. At the same time, people tend to overestimate the effectiveness of charities that they find personally appealing [21], especially ones that are local [27,71]. Many people also assume that charities responding to disasters are especially effective, although charities addressing ongoing problems in poor countries (e.g., neglected tropical diseases) tend to be more effective [21,72]. This may be driven by people's greater emotional responses to victims who have experienced distinct losses compared with victims of chronic misfortune [73].

Strategies to Increase Effective Giving

Previous research has focused on strategies to increase people's willingness to give to charity and the amounts that people give [4], but there is relatively little research on strategies encouraging people to give more effectively. Whether or not it is good to give more effectively, or to give at all, is a moral question [5–7] (Box 2) that we do not attempt to answer here. Nevertheless, we proceed on the assumption that saving and improving more lives rather than fewer lives is a good thing, and discuss psychologically informed strategies to encourage more effective giving.

Information

A natural strategy is to inform donors about which charities are most effective (as GiveWell does) and dispel misconceptions about effectiveness (e.g., the Overhead Myth Campaign^v). The evidence for the effectiveness of effectiveness information is mixed. Some studies have found no effect [74,75], while some of our own studies have found a large effect, increasing the proportion of effective donors from zero to 17% or even higher [21,76]. Other studies point to heterogeneity among donors, with minimal effects overall but larger effects for donors who are more altruistically motivated [77] and more educated, especially when encouraged to think more deliberately [78]. Providing more tangible details about a charity's intervention strategies may also make giving more effective [75,79]. All studies to date use limited amounts of effectiveness information, but some donors, including donors willing to give large amounts, may be influenced by more extensive information.

Choice Architectures and Incentives

Defaults

As with other behaviors [80], setting donating as the default behavior can increase its prevalence [81]. Default settings could also be used to increase the effectiveness of giving. While default choices are by no means specific to effective giving, highly effective defaults are especially relevant because they can be justified in terms of their greater social benefits, as in the classic case of opt-out organ donation [82].

Incentivizing

Another general strategy to promote a behavior is to provide incentives. Consistent with this, charities benefit from donor tax deductions and often use matching campaigns [83]. Although incentivization is a general strategy, it may be especially relevant to effective giving if there is a subset of donors who are highly motivated to incentivize others to give effectively and who are therefore willing to provide matching funds for this purpose (see later regarding Giving Multiplier). Likewise, governments could provide larger incentives for effective giving through targeted tax deductions.

Unit Asking

Scope neglect (see earlier) is reduced when people evaluate donation options jointly rather than separately [41]. The 'unit asking' technique builds on this idea [84]. Potential donors are

first asked how much they would give to help one person in need before being asked how much they would give to help a group of a specified size. Unit asking tends to increase donors' inclinations to raise their donations proportionally with group size and thus makes giving more effective.

Splitting

As noted earlier, people tend to split donations across multiple charities when given the option [49,85]. While the preference for splitting can reduce the effectiveness of giving (see earlier), it can also promote more effective giving if the tendency to split causes people to allot some, rather than none, of their money to a highly effective charity. As discussed earlier, people tend to support charities that are emotionally appealing, paying little attention to effectiveness. However, there is evidence that many people do care about effectiveness and that information about effectiveness can make giving more effective [2,21]. Combining these insights suggests a new strategy to increase the effectiveness of charitable giving: many donors may be amenable to splitting their donations between an emotionally appealing charity and a highly effective charity, especially if provided with effectiveness information. We have applied this technique in the design of an online donation platform (Giving Multiplier^{vi}), which uses matching funds (see earlier) to encourage donors to split their donations between a highly effective charity and their favorite charity [76].

Philosophical Reasoning

Philosophical arguments of the kind pioneered by Peter Singer [86] (Box 2) can increase people's willingness to donate to an effective charity [87,88]. Likewise, veil-of-ignorance reasoning of the kind pioneered by Rawls [89] and Harsanyi [90] – asking people to consider what they would want if they did not know who they are going to be – can counteract parochial tendencies ('narrow moral circle') and thus make giving more effective [25]. The effective altruism movement was inspired by philosophers [1,86,91,92] and has directed billions of dollars toward effective charities [93], indicating that philosophical arguments can have a powerful impact (see also Open Philanthropy^{vii}). However, it may be that relatively few people are receptive to such arguments [87,88]. Understanding how philosophical arguments (Box 2) might influence the behavior of more typical donors is an important topic for future research.

Box 2. Is Effectiveness Morally Obligatory?

In one of the most influential philosophical arguments of the 20th century, Peter Singer asked readers to imagine encountering a child drowning in a shallow pond [86]. Is it morally acceptable to allow the child to drown because one does not want to ruin one's clothes? Singer reasonably assumed that it would be wrong to allow the child to drown. However, he argued further that affluent people have a comparable moral obligation to save the lives of distant children, at least when relatively little personal sacrifice is required. While many have found this argument convincing, many more have found it unsettling: it implies that even moderately affluent people ought to give far more to charity than most of them currently give.

More recently, philosophers have considered a related moral challenge that applies even to those who reject Singer's conclusion: if one is going to donate, does one have an obligation to choose one of the most effective charities? Parfit considered a case in which a man is about to lose both of his arms [107]. Someone can save either one arm or both arms at great cost to themselves, but the cost is the same whether they save one arm or both. Parfit grants that one is not obliged to save a stranger's limb(s) at great cost to oneself. However, he claims that, if one decides to help, it would be perverse not to save both arms.

Applying this logic to charitable giving, Pummer argues that even if it is not obligatory to donate to charity, it is obligatory to donate effectively once one has decided to give [108]. MacAskill [109] considers an interpersonal version of Parfit's case. Suppose that a firefighter who is charging into a burning building has time to reach only one room. One room contains one child, whereas the other contains ten. The conclusion, says MacAskill, is the same: If one is going to help, it is wrong to help in a way that helps less effectively.

In the firefighter case, the more effective option is ten times more effective. Ord's analysis of global health interventions indicates that this ratio undersells the moral imperative to give effectively [1]. He finds that the most effective interventions can be hundreds or even thousands of times more effective than other interventions (Figure 1).

Norm Changes

One of the most robust findings in research on social influence is the power of norms [94]. As explained earlier, prevailing altruistic norms emphasize self-sacrifice, but not effectiveness. There is also a norm against publicizing one's altruistic behavior [95,96] and (outside certain religions) there is little expectation that people donate a significant percentage of their disposable income. In the effective altruism community, these norms have shifted [5,6]. Many effective altruists publicly commit to significant life-long giving to effective charities in the hope of encouraging others to do the same (see Giving What We Can^{viii}, The Life You Can Save^{ix}). There is also a visible commitment to evidence-based decision making and open-mindedness to new causes. It is unknown how widely such effectiveness-oriented norms could spread, but history tells us that radical norm change is possible, as demonstrated by widespread changes in moral views about slavery and racial discrimination, women's rights, gay rights, etc. [97,98].

Concluding Remarks

Ordinary people have the power to save and transform people's lives through effective giving, and we are beginning to understand the factors that encourage and discourage such choices, with many open questions for future research (see Outstanding Questions).

Some obstacles to effective giving are motivational: People believe that whether and how to give is primarily a matter of personal preference. Altruism tends to be motivated by narrow affective responses that are highly sensitive to personal relevance and insensitive to the scope of the problems that may be addressed. People are also averse to prioritizing some causes over others and receive little reputational benefit (and may even pay reputational costs) for prioritizing effectiveness.

Other obstacles are epistemic: people mistakenly believe that helping indirectly through organizational overhead is wasteful and that effectiveness cannot be well quantified. People misunderstand probability, fail to think in terms of marginal value, and dismiss important benefits as a drop in the bucket. Furthermore, people are unaware that charities vary enormously in their effectiveness, and do not know which charities are most effective.

There are, however, promising strategies to address these obstacles to effective giving. More accurate information can remove epistemic obstacles. Well-designed choice architectures and incentives may overcome motivational limitations, boosting people's existing (but rarely dominant) concerns for effectiveness. Finally, in the long run, shifting norms may lead to widespread commitments to doing as much good as possible with our resources.

Acknowledgments

We thank Samantha Kassirer, Joshua Lewis, John Halstead and Matt Coleman for their valuable inputs on the manuscript. This research was funded by EA Funds (L.C.) and EA Grants (S.S.).

Declaration of Interests

The authors have no interests to declare.

Resources

ⁱ<http://effectivealtruism.org>

ⁱⁱ<https://animalcharityevaluators.org/>

ⁱⁱⁱ<https://80000hours.org/>

^{iv}<http://givewell.com/>

^v<http://overheadmyth.com/>

Outstanding Questions

What motivates effective giving? We have focused on obstacles to effective giving but have said little about positive motivations for effectiveness. On a cognitive level, what processes support effective giving? At the level of individual differences, what traits, beliefs, and values motivate effective giving? Existing studies show that a sizable minority of people are moved to give effectively when informed about charity effectiveness. What are their distinguishing traits and how are they acquired?

In what ways do attitudes toward effectiveness in charitable giving (and altruism more generally) vary across cultures? Beyond current attitudes, how might cultural groups vary in their amenability to effective giving? At present, effective altruism has drawn adherents primarily from Western societies, but its underlying philosophy has parallels in non-Western cultures and it aspires to maximize global benefits. Is effective giving inherently culturally bound or is it more cross-culturally amenable?

How has effectiveness been valued throughout history and across domains (e.g., in medicine)? What lessons can be learned from prior attempts to implement novel evidence-based practices?

What interventions can promote effective giving? Information campaigns and behavioral interventions should be tested in field studies. Researchers might also examine long-term interventions aimed at changing norms.

- vi <http://givingmultiplier.org/>
- vii <https://openphilanthropy.org/>
- viii <https://givingwhatwecan.org/>
- ix <http://thelifeyoucansave.org>

References

1. Ord, T. (2013) *The Moral Imperative Toward Cost-Effectiveness in Global Health*, Center for Global Development
2. Caviola, L. et al. (2020) Donors vastly underestimate differences in charities' effectiveness. *Judgm. Decis. Mak.* 15, 509–516
3. Giving USA Foundation (2020) *Giving USA 2020: The Annual Report on Philanthropy for the Year 2019*, Giving USA
4. Bekkers, R. and Wiepking, P. (2011) A literature review of empirical studies of philanthropy: eight mechanisms that drive charitable giving. *Nonprofit Volunt. Sect. Q.* 40, 924–973
5. MacAskill, W. (2015) *Doing Good Better: Effective Altruism and a Radical New Way to Make a Difference*, Guardian Faber
6. Singer, P. (2015) *The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically*, Yale University Press
7. Fiennes, C. (2017) We need a science of philanthropy. *Nature* 546, 187
8. Batson, D. (2014) *The Altruism Question: Toward A Social-Psychological Answer*, Psychology Press
9. Berman, J.Z. et al. (2018) Impediments to effective altruism: the role of subjective preferences in charitable giving. *Psychol. Sci.* 29, 834–844
10. Camber Collective (2015) *Money for Good 2015: Revealing the Voice of the Donor in Philanthropic Giving*, Camber Collective
11. Collinger, T. (2017) *How Online Reviews Influence Sales*, Medill I/MC Spiegel Research Center
12. Lerner, J.S. and Tetlock, P.E. (1999) Accounting for the effects of accountability. *Psychol. Bull.* 125, 255
13. Nagel, J. and Waldmann, M.R. (2013) Deconfounding distance effects in judgments of moral obligation. *J. Exp. Psychol. Learn. Mem. Cogn.* 39, 237–252
14. Caviola, L. and Schubert, S. (2020) Is it obligatory to donate effectively? Judgments about the wrongness of donating ineffectively. *PsyArXiv* Published online August 24, 2020. <https://doi.org/10.31234/osf.io/2h4r>
15. Batson, C.D. et al. (1981) Is empathic emotion a source of altruistic motivation? *J. Pers. Soc. Psychol.* 40, 290–302
16. Batson, C.D. et al. (1987) Distress and empathy: two qualitatively distinct vicarious emotions with different motivational consequences. *J. Pers.* 55, 19–39
17. Loewenstein, G. and Small, D.A. (2007) The Scarecrow and the Tin Man: the vicissitudes of human sympathy and caring. *Rev. Gen. Psychol.* 11, 112–126
18. Andreoni, J. (1990) Impure altruism and donations to public goods: a theory of warm-glow giving. *Econ. J. Nepal* 100, 464–477
19. Bloom, P. (2016) *Against Empathy*, HarperCollins
20. Cushman, F. (2013) Action, outcome, and value a dual-system framework for morality. *Personal. Soc. Psychol. Rev.* 17, 273–292
21. Caviola, L. et al. (2020) The many obstacles to effective giving. *Judgm. Decis. Mak.* 15, 159
22. Small, D.A. and Simonsohn, U. (2007) Friends of victims: personal experience and prosocial behavior. *J. Consum. Res.* 35, 532–542
23. Samuelson, W. and Zeckhauser, R. (1988) Status quo bias in decision making. *J. Risk Uncertain.* 1, 7–59
24. Jones, B. and Rachlin, H. (2006) Social discounting. *Psychol. Sci.* 17, 283–286
25. Huang, K. et al. (2019) Veil-of-ignorance reasoning favors the greater good. *Proc. Natl. Acad. Sci. U. S. A.* 116, 23989–23995
26. Baron, J. and Szymanska, E. (2011) Heuristics and biases in charity. In *The Science of Giving: Experimental Approaches to the Study of Charity* (Oppenheimer, D.M. and Olivola, C.Y., eds), pp. 215–235, Psychology Press
27. Knowles, S. and Sullivan, T. (2017) Does charity begin at home or overseas? *Nonprofit Volunt. Sect. Q.* 46, 944–962
28. Caviola, L. et al. (2019) The moral standing of animals: towards a psychology of speciesism. *J. Pers. Soc. Psychol.* 116, 1011–1029
29. Caviola, L. et al. (2020) Utilitarianism for animals, Kantianism for people? Harming animals and humans for the greater good. *J. Exp. Psychol. Gen.* Published online October 19, 2020. <https://doi.org/10.1037/xge0000988>
30. Ord, T. (2020) *The Precipice: Existential Risk and the Future of Humanity*, Hachette
31. Schubert, S. et al. (2019) The psychology of existential risk: moral judgments about human extinction. *Sci. Rep.* 9, 15100
32. Dickert, S. et al. (2015) Scope insensitivity: the limits of intuitive valuation of human lives in public policy. *J. Appl. Res. Mem. Cogn.* 4, 248–255
33. Fetherstonhaugh, D. et al. (1997) Insensitivity to the value of human life: a study of psychophysical numbing. *J. Risk Uncertain.* 14, 283–300
34. Västfjäll, D. et al. (2014) Compassion fade: affect and charity are greatest for a single child in need. *PLoS One* 9, e100115
35. Slovic, P. (2007) When compassion fails. *New Scientist* 194, 18
36. Desvousges, W.H. et al. (1992) *Measuring Nonuse Damages Using Contingent Valuation: An Experimental Evaluation of Accuracy*, Research Triangle Institute
37. Kahneman, D. and Knetsch, J.L. (1992) Valuing public goods: the purchase of moral satisfaction. *J. Environ. Econ. Manag.* 22, 57–70
38. Shenhav, A. and Greene, J.D. (2010) Moral judgments recruit domain-general valuation mechanisms to integrate representations of probability and magnitude. *Neuron* 67, 667–677
39. Slovic, P. (2010) If I look at the mass I will never act: psychic numbing and genocide. In *Emotions and Risky Technologies*, pp. 37–59, Springer
40. Kogut, T. and Ritov, I. (2005) The "identified victim" effect: an identified group, or just a single individual? *J. Behav. Decis. Mak.* 18, 157–167
41. Kogut, T. and Ritov, I. (2005) The singularity effect of identified victims in separate and joint evaluations. *Organ. Behav. Hum. Decis. Process.* 97, 106–116
42. Sudhir, K. et al. (2016) Do sympathy biases induce charitable giving? The effects of advertising content. *Mark. Sci.* 35, 849–869
43. Small, D.A. and Loewenstein, G. (2003) Helping a victim or helping the victim: altruism and identifiability. *J. Risk Uncertain.* 26, 5–16
44. Bloom, P. (2013) *The baby in the well: the case against empathy*, *New Yorker*, May 20
45. Small, D.A. et al. (2007) Sympathy and callousness: the impact of deliberative thought on donations to identifiable and statistical victims. *Organ. Behav. Hum. Decis. Process.* 102, 143–153
46. Tetlock, P.E. (2003) Thinking the unthinkable: sacred values and taboo cognitions. *Trends Cogn. Sci.* 7, 320–324
47. Ubel, P.A. et al. (1996) Cost-effectiveness analysis in a setting of budget constraints – is it equitable? *N. Engl. J. Med.* 334, 1174–1177
48. Paolacci, G. and Yalcin, G. (2020) Fewer but poorer: benevolent partiality in prosocial preferences. *Judgm. Decis. Mak.* 15, 173–181
49. Sharps, D.L. and Schroeder, J. (2019) The preference for distributed helping. *J. Pers. Soc. Psychol.* 117, 954–977
50. Null, C. (2011) Warm glow, information, and inefficient charitable giving. *J. Public Econ.* 95, 455–465
51. Exley, C.L. (2020) Using charity performance metrics as an excuse not to give. *Manag. Sci.* 66, 553–563

52. Gneezy, U. et al. (2014) Avoiding overhead aversion in charity. *Science* 346, 632–635
53. Cunningham, H. et al. (2017) Bilateral foreign aid: how important is aid effectiveness to people for choosing countries to support? *Appl. Econ. Lett.* 24, 306–310
54. Uhlmann, E.L. et al. (2015) A person-centered approach to moral judgment. *Perspect. Psychol. Sci.* 10, 72–81
55. Nowak, M.A. (2006) Five rules for the evolution of cooperation. *Science* 314, 1560–1563
56. Yoeli, E. et al. (2013) Powering up with indirect reciprocity in a large-scale field experiment. *Proc. Natl. Acad. Sci. U. S. A.* 110, 10424–10429
57. Burum, B. et al. (2020) An evolutionary explanation for ineffective altruism. *Nat. Hum. Behav.* 4, 1245–1257
58. Johnson, S. (2018) Dimensions of altruism: do evaluations of prosocial behavior track social good or personal sacrifice? *SSRN* Published online November 15, 2018. <https://doi.org/10.2139/ssrn.3277444>
59. Barasch, A. et al. (2014) Selfish or selfless? On the signal value of emotion in altruistic behavior. *J. Pers. Soc. Psychol.* 107, 393–413
60. Simler, K. and Hanson, R. (2017) *The Elephant in the Brain: Hidden Motives in Everyday Life*, Oxford University Press
61. Montealegre, A. et al. (2020) Does maximizing good make people look bad? *PsyArXiv* Published online April 12, 2020. <https://doi.org/10.31234/osf.io/2zbx>
62. Caviola, L. et al. (2014) The evaluability bias in charitable giving: saving administration costs or saving lives? *Judgm. Decis. Mak.* 9, 303
63. Torgovnick May, K. (2014) *Charities + Overhead: Dan Pallotta's TED Talk Changes the Conversation*, TED
64. Berrett, J.L. (2020) *Burying the Overhead Myth and Breaking the Nonprofit Starvation Cycle: Identifying More Valid Measures and Determinants of Nonprofit Efficiency*, North Carolina State University
65. Zeckhauser, R. and Shepard, D. (1976) Where now for saving lives. *Law Contemp. Probs* 40, 5
66. Banerjee, A.V. et al. (2011) *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*, PublicAffairs
67. Peters, E. et al. (2006) Numeracy and decision making. *Psychol. Sci.* 17, 407–413
68. Snowden, J. (2019) Should we give to more than one charity? In *Effective Altruism: Philosophical Issues* (Greaves, H. and Pummer, T., eds), Oxford University Press
69. Baron, J. and Greene, J. (1996) Determinants of insensitivity to quantity in valuation of public goods: contribution, warm glow, budget constraints, availability, and prominence. *J. Exp. Psychol. Appl.* 2, 107
70. Västfjäll, D. et al. (2015) Pseudoinefficacy: negative feelings from children who cannot be helped reduce warm glow for children who can be helped. *Front. Psychol.* 6, 616
71. Grimson, D. et al. (2020) How close to home does charity begin? *Appl. Econ.* 52, 3700–3708
72. Karnofsky, H. (2013) *6 Tips on Disaster Relief Giving*, GiveWell
73. Small, D.A. (2010) Reference-dependent sympathy. *Organ. Behav. Hum. Decis. Process.* 112, 151–160
74. Metzger, L. and Günther, I. (2019) Making an impact? The relevance of information on aid effectiveness for charitable giving. A laboratory experiment. *J. Dev. Econ.* 136, 18–33
75. Bergh, R. and Reinstein, D. (2020) Empathic and numerate giving: the joint effects of victim images and charity evaluations. *Soc. Psychol. Pers. Sci.* Published online February 17, 2020. <https://doi.org/10.1177/1948550619893968>
76. Caviola, L. and Greene, J.D. (2021) Boosting effective giving with bundling and donor coordination. *PsyArXiv* Published online March 22, 2021. <https://doi.org/10.31234/osf.io/65fmr>
77. Karlan, D. and Wood, D.H. (2017) The effect of effectiveness: donor response to aid effectiveness in a direct mail fundraising experiment. *J. Behav. Exp. Econ.* 66, 1–8
78. Verkaik, D. (2016) Do donors really care about impact information? A dual process account. *Open Sci. Framework* Published online May 18, 2016. <https://osf.io/paf2b>
79. Cryder, C.E. et al. (2013) The donor is in the details. *Organ. Behav. Hum. Decis. Process.* 120, 15–23
80. Thaler, R.H. and Sunstein, C.R. (2009) *Nudge*, Penguin
81. Everett, J.A.C. et al. (2015) Doing good by doing nothing? The role of social norms in explaining default effects in altruistic contexts. *Eur. J. Soc. Psychol.* 45, 230–241
82. Johnson, E.J. and Goldstein, D. (2003) Do defaults save lives? *Science* 302, 1338–1339
83. Karlan, D. and List, J.A. (2007) Does price matter in charitable giving? Evidence from a large-scale natural field experiment. *Am. Econ. Rev.* 97, 1774–1793
84. Hsee, C.K. et al. (2013) Unit asking: a method to boost donations and beyond. *Psychol. Sci.* 24, 1801–1808
85. Weisz, E. and Cikara, M. (2020) Merely increasing action options increases charitable donation. *Research Square* Published online August 28, 2020. <https://doi.org/10.21203/rs.3.rs-59021/v1>
86. Singer, P. (1972) Famine, affluence, and morality. *Philos Public Aff* 1, 229–243
87. Lindauer, M. et al. (2020) Comparing the effect of rational and emotional appeals on donation behavior. *Judgm. Decis. Mak.* 15, 413–420
88. Schwitzgebel, E. and Cushman, F. (2020) Contest winner! A philosophical argument that effectively convinces research participants to donate to charity. *Splintered Mind* Published online June 23, 2020. <http://schwitzsplinters.blogspot.com/2020/06/contest-winner-philosophical-argument.html>
89. Rawls, J. (2009) *A Theory of Justice*, Harvard University Press
90. Harsanyi, J.C. (1955) Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *J. Polit. Econ.* 63, 309–321
91. Singer, P. (2016) *The Most Good You Can Do*, Yale University Press
92. MacAskill, W. (2015) *Doing Good Better*, Guardian Faber
93. MacAskill, W. (2019) The definition of effective altruism. In *Effective Altruism: Philosophical Issues* (Greaves, H. and Pummer, T., eds), Oxford University Press
94. Cialdini, R.B. and Goldstein, N.J. (2004) Social influence: compliance and conformity. *Annu. Rev. Psychol.* 55, 591–621
95. Berman, J.Z. et al. (2015) The braggart's dilemma: on the social rewards and penalties of advertising prosocial behavior. *J. Mark. Res.* 52, 90–104
96. De Freitas, J. et al. (2019) Maimonides' ladder: states of mutual knowledge and the perception of charity. *J. Exp. Psychol. Gen.* 148, 158–173
97. Pinker, S. (2012) *The Better Angels of Our Nature: Why Violence Has Declined*, Penguin
98. Singer, P. (2011) *The Expanding Circle: Ethics, Evolution, and Moral Progress*, Princeton University Press
99. Johnson, R. et al. (2016) Where's WALY? A proof of concept study of the "wellbeing adjusted life year" using secondary analysis of cross-sectional survey data. *Health Qual. Life Outcomes* 14, 126
100. Miguel, E. and Kremer, M. (2004) Worms: identifying impacts on education and health in the presence of treatment externalities. *Econometrica* 72, 159–217
101. Baird, S. et al. (2016) Worms at work: long-run impacts of a child health investment. *Q. J. Econ.* 131, 1637–1680
102. Mintz-Woo, K. (2020) A philosophers' guide to discounting. In *Philosophy & Climate Change* (McPherson, T. et al., eds), Oxford University Press
103. Stern, N.H. et al. (2006) *Stern Review on the Economics of Climate Change*, Cambridge University Press
104. Singer, P. (1975) *Animal Liberation: A New Ethic for Our Treatment of Animals*, Avon
105. Greaves, H. (2017) Population axiology. *Philos Compass* 12, e12442
106. Greaves, H. et al. (2020) *A Research Agenda for the Global Priorities Institute*, University of Oxford
107. Parfit, D. (1986) *Reasons and Persons*, Oxford University Press
108. Pummer, T. (2016) Whether and where to give. *Philos Public Aff* 44, 77–95

109. MacAskill, W. (2013) What Charity Navigator Gets Wrong About Effective Altruism. *Stanf. Soc. Innov. Rev.* December 3
110. Jamison, D.T. *et al.* (2006) *Disease Control Priorities in Developing Countries*, World Bank
111. Layard, R. and Oparina, E. (2021) *Living Long and Living Well: The WELLBY Approach*, World Happiness Report, p. 191
112. Hamory, J. *et al.* (2021) Twenty-year economic impacts of deworming. *Proc. Natl. Acad. Sci.* 118, e2023185118