

Effective Altruism

1. Introduction

There are many problems in the world today. 750 million people live on less than \$2.00 (PPP adjusted) per day.¹ 7 million children die each year of easily-preventable causes such as malaria, diarrhea or pneumonia.² Climate change is set to wreak environmental havoc and cost the economy trillions of dollars.³ A third of women worldwide have suffered from sexual or other physical violence in their lives.⁴ Bacteria are becoming antibiotic-resistant. 2500 nuclear warheads are in high-alert ready-to-launch status around the globe.⁵

Given that the world has so many problems, and that these problems are so severe, surely we have a responsibility to do something to leave the world a better place than how we found it. But what? There are millions of problems that we could be addressing, and thousands of ways of addressing each of those problems. And our resources are scarce, so as individuals and even as a globe we can't solve all these problems at once. So we must make decisions about how to allocate the resources we have. But on what basis should we make such decisions?

The effective altruism movement has pioneered one approach. Those in this movement try to figure out, of all the different uses of our resources, which uses will do the most good, impartially considered. This movement is gathering considerable steam, with over 3000 people taking *Giving What We Can*'s pledge to give at least 10% of their income for the rest of their lives to the organisations they believe to be most cost-effective, \$90 million per year moved to GiveWell's top recommended charities, a foundation with potential assets of \$10 billion joining the fold, and a community consisting of thousands of people around the world who are trying to use their time on earth to improve the world by as much as they can.

As defined by the leaders of the movement, effective altruism is the use of evidence and reason to work out how to benefit others by as much as possible, and the taking action on that basis. So defined, effective altruism is a project, rather than a set of normative commitments. It is both a

¹ World Bank Group, *Poverty and Shared Prosperity: Taking on Inequality 2016* (Herndon: World Bank Publications, 2016), ch. 2

² <https://www.unicef.org/mdg/childmortality.html>, UNICEF, "Levels & Trends in Child Mortality", https://www.unicef.org/publications/files/Child_Mortality_Report_2017.pdf.

³ John Broome, *Climate Matters: Ethics in a Warming World* (New York: W.W. Norton, 2012); William Nordhaus, *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World* (New Haven: Yale University Press, 2015).

⁴ United Nations, *The World's Women 2015: Trends and Statistics* (New York: United Nations, 2015), https://unstats.un.org/unsd/gender/downloads/worldswomen2015_report.pdf

⁵ https://en.wikipedia.org/wiki/High-alert_nuclear_weapon

research project — to figure out how to do the most good — and a practical project, of implementing the best guesses we have about how to do the most good.

However, as moral philosophers we can still ask whether effective altruism is a project that we should be pursuing. In this article, I argue in favour of effective altruism's approach. More precisely, I argue (i) that those of us who are well off have an obligation to make helping others a significant part of our lives, and (ii) that when helping others, we have an obligation to do so in the most effective way. I will begin in the next section by arguing in favour of the first claim, which I call *Duty of Beneficence*. After that I argue for the second claim, which I call *Maximising Beneficence*. Finally, I will outline a framework for figuring out which ways of helping others are most effective.

2. Duty of Beneficence

My first claim is that those of us who are well off have a significant obligation to help others. More specifically, I will argue for:

Duty of Beneficence: Most middle or upper class people in rich countries have a duty to make helping others a significant part of their lives.

Although I won't give a precise definition of what it means to make helping others a "significant" part of one's life, here are some examples to give you an idea of what I have in mind:

- Give 10% of one's income to charity
- Let social impact be a major consideration when choosing a career/industry/job
- Engaging in volunteering/community organising/advocacy for 10hrs/week

Arguments for a Duty of Beneficence

Why should we think there is such a duty of beneficence?

First of all, most moral views recognise some duty of beneficence. Let us consider three influential schools of Western moral philosophy: utilitarianism, virtue ethics, and Kantian deontology. The implication is perhaps most straightforward in the case of utilitarianism. If you belong to the middle class in a rich country, then the well-being that you forego by donating, say, \$1000 is trivial compared to the amount of good that money could do to benefit others in the developing world. Instead of buying a new smartphone, your donation could give someone the

equivalent of several years of healthy life. As the latter would clearly lead to more happiness in the world, utilitarianism commands you to do it.

In contrast with utilitarianism, the fundamental ethical question for virtue ethics is: “What sort of person should I be?” And the answer is that one should be a virtuous person, namely one who possesses and exercises virtuous character traits such as courage, honesty, justice, generosity, etc. Plausibly, the appropriate form of generosity entails making helping others a part of one’s life, and virtue ethics therefore also recognises a duty of beneficence.

Finally, according to one formulation of Kant’s categorical imperative, an act is only permissible if it is based on a maxim that we could rationally will to be universal law. Recognising that a maxim of never helping others fails this “universalizability” test, Kant claims that “[w]e have a duty to be charitably helpful where we can.”⁶ When so many different moral views are in agreement, that gives us good reason to think that there is some duty of beneficence.

But does a duty of beneficence require a ‘significant’ commitment? There are three reasons for thinking that it does.

First, the sheer amount of good that one can do with one’s resources is huge. According to the best-guess estimate from GiveWell, an organisation that conducts exceptionally in-depth charity evaluations, it costs about \$3500 to save a life.⁷ By donating 10% of their income each year, a middle-class member of a rich country will save a child’s life every couple of years — dozens of lives over their lifetime. And, if that person were to focus on other cause-areas (discussed in section 4), it’s plausible that they could do far more good again.

Second, middle-class members of rich countries are exceptionally wealthy by global standards. The richest 3% of people worldwide earn 50 times as much as the poorest 750 million.⁸ Given this incredible disparity of riches, surely that global elite has an obligation to use at least a significant proportion of their wealth to help the poor and disenfranchised. But what’s not intuitive, though it’s true, is that most middle class members of rich countries are among the richest 3% of the world’s population. So if you think the global elite has an obligation to help, and you are in the richest 3% worldwide,⁹ then you should make helping others a significant part of your life.

⁶ Groundwork, Chapter 1, transl. by Jonathan Bennett, p. 8 here:
<http://www.earlymoderntexts.com/assets/pdfs/kant1785chapter1.pdf>

⁷ <https://www.givewell.org/how-we-work/our-criteria/cost-effectiveness/cost-effectiveness-models>

⁸ <https://www.givingwhatwecan.org/post/2015/12/do-something-incredible-this-new-year/>

⁹ Find out here: <https://www.givingwhatwecan.org/get-involved/how-rich-am-i/>

Third, significant personal commitments involve sacrificing far less than one might initially have thought. Let's consider charitable giving and career choice in turn.

Studies suggest that although there is a positive correlation between income and happiness, it is not as strong as one might think. In the US, for example, a 10% reduction in income is associated with only a 1% drop on a scale measuring life satisfaction.¹⁰ Moreover, it is not at all clear that we should think of donating 10% as equivalent to a 10% loss of income. There is evidence to suggest that spending money on others can often improve our well-being more than spending money on ourselves.¹¹ So it's not even clear that donating 10% of one's income would be much of a personal sacrifice at all.¹² Similarly, engaging in volunteer work, far from being a personal sacrifice, is associated with improved well-being along a wide variety of dimensions.¹³

Secondly, in the case of career choice, we are able to enjoy a much broader variety of jobs than we often think before we've tried them.¹⁴ What's more, you are unlikely to thrive in a job that you don't enjoy. Hence it would be unsustainable to suggest that you pursue a career doing something that you hate. Therefore, in letting social impact be a major consideration, you would not be giving up on a satisfactory, challenging, and enjoyable career. Thus, for both charitable giving and career choice, the Duty of Beneficence turns out to be less demanding than one might think.

3. Maximising Beneficence

If, as I have just argued, those of us who are well off have a duty to make helping others a significant part of our lives, how should we go about performing this duty? How do we choose among the many different ways of helping others? I will argue that we should help others in the most *effective* way possible. More specifically, I will argue for:

Maximising Beneficence: With respect to those resources that we have a duty of beneficence to use to improve the world, and subject to not violating anyone's rights, it is imperative that we try to use our resources to do the most good, impartially considered, that we can.

The case for Maximising Beneficence is simple. If we fail to produce the best outcome we can, that means that more people die than needed to die, or that people suffer larger harms than they

¹⁰ Stevenson and Wolfers (2013). Subjective Well-Being and Income: Is There Any Evidence of Satiation? <https://www.nber.org/papers/w18992>

¹¹ Dunn, Gilbert & Wilson (2011). If money doesn't make you happy, then you probably aren't spending it right. <https://psycnet.apa.org/record/2011-17293-002>

¹² For more detail, see MacAskill, Mogensen, and Ord (2018). Giving Isn't Demanding

¹³ Thoits and Hewitt (2001). Volunteer Work and Well-Being

¹⁴ <https://80000hours.org/career-guide/job-satisfaction/>

needed to suffer, or that the better off are benefitted at the expense of the worse off. If we think that the grave harms that others in this world suffer are urgent enough that we ought to use some of our resources to fight those harms, those exact same reasons warrant us using those resources in ways that help as much as possible.

The practical significance of Maximising Beneficence becomes clear once we realise that different ways of doing good differ vastly in the amount of good that they can do. This is counterintuitive: on average, people think that the best charities differ in their effectiveness compared with typical charities only by a factor of 1.5 or so.¹⁵ However, the most cost-effective charities are far more effective than typical charities — plausibly a hundred times more.

Even in the developing world, different programs — each of which we would consider a good use of money — vary by tens or hundreds-fold. According to the independent charity evaluator GiveWell, donations to the Against Malaria Foundation will provide a benefit equivalent to a year of healthy life (or QALY¹⁶) for about \$100.¹⁷ In contrast, the cost to provide one QALY on the margin via the NHS is \$40,000.¹⁸ This means that donations to Against Malaria Foundation are giving 400 times as many years of healthy life as UK medical spending is. Insofar as medical non-profits in the UK will be operating beyond this margin (and in general don't have as strict policies on ensuring cost-effectiveness), we should expect their cost-effectiveness to be lower than this. Once we move to cost per life saved, the difference gets even more stark. According to GiveWell's best estimate, it costs around \$3500 to save a life via distributing bednets.¹⁹ The typical cost spent by the US government to save a life is more like \$9.5 million.²⁰

This isn't merely limited to healthcare charities. The same phenomenon is true among education charities. Among estimates from the Poverty Action Lab, the same phenomenon holds: the difference in improvements of test scores with a given amount of money between the most effective programs and typical programs (which would be regarded as highly cost-effective in a domestic setting) is a factor measured in the hundreds.²¹

In light of these vast differences between charities, the decision of where to donate is of huge consequence, and Maximising Beneficence requires us to make this decision very carefully. The

¹⁵

<https://80000hours.org/2017/05/most-people-report-believing-its-incredibly-cheap-to-save-lives-in-the-developing-world/>

¹⁶ Quality-adjusted life years. See https://en.wikipedia.org/wiki/Quality-adjusted_life_year

¹⁷ <https://blog.givewell.org/2016/12/12/amf-population-ethics/>

¹⁸ <https://www.nice.org.uk/process/pmg9/chapter/the-appraisal-of-the-evidence-and-structured-decision-making>

¹⁹ <https://www.givewell.org/how-we-work/our-criteria/cost-effectiveness/cost-effectiveness-models>

²⁰ <https://www.bloomberg.com/graphics/2017-value-of-life/>

²¹ See <https://www.povertyactionlab.org/policy-lessons/education/increasing-test-score-performance>. Note, however, that these estimates have not been subject to the same level of scrutiny as GiveWell's estimates.

most striking way this differs from prevailing attitudes to doing good is with respect to choice of causes. An implication of Maximising Beneficence is what we may call *cause-impartiality*: the choice of cause that one focuses on should be determined only by the amount of good that one can do by focusing on that cause. In contrast, the prevailing view with respect to doing good is that one's choice of cause is a matter of personal preference. Perhaps, once one has chosen a cause, one should try to do as much good as possible with respect to that cause. But whether one should focus on US education, or the arts, or climate change — there's no objective answer to that question, and the best approach is to work on whatever cause you're most personally passionate about.

However, I do not think that personal relationships to cause-areas provides a good response to *Maximising Beneficence*. Two arguments support this view. First is simply an appeal to cases. Suppose that, as a volunteer doctor in a resource-starved hospital in a poor country, you can do one of two things with your last day of work before you return home. First, you could perform surgery on an elderly man with prostate cancer, thereby saving his life. Or you could treat two children from malaria, thereby saving both their lives. If you had a personal attachment to the cause of fighting prostate cancer, would that give you sufficient reason to save the the life of the elderly man rather than the two children? Clearly not. The importance of saving two lives rather than one, and of saving people who have much more to gain from their treatment, clearly outweighs whatever reason a personal attachment might bring. Yet this is morally analogous to the decisions that we actually face when we try to use our resources to do good. The only way in which it is morally disanalogous is with respect to what's at stake. In the world we live in, the choice typically isn't between preventing the death of one person from prostate cancer and two children from malaria; it's between preventing the death of one person from prostate cancer and preventing the deaths of dozens of children from malaria.

The second argument is based on the arbitrariness of any way of generalising from a personal attachment to an individual. Suppose, again, that I had a family member who died of prostate cancer. Does that give me an additional reason to support charities that work on prostate cancer? Why would it give me a reason to support charities that fight prostate cancer, rather than all cancer in general, or all non-transmissible diseases? Or all forms of suffering? And if there was some compelling argument that moved from my personal circumstances to a reason in favour of supporting charities that fight prostate cancer, why should that argument not support a narrower problem than that? Why should I not support the cause of British people with prostate cancer? Or people who are over 70 and have prostate cancer? In general, there is no non-arbitrary way of delineating causes. And that means that any argument from personal circumstances to partiality among causes will have to pick from one arbitrary delineation over another.

On the basis of these two arguments, it seems to me that attachments to causes cannot serve as an argument against the idea that we should do the most good with the resources we have a duty to use to improve the world. This is not to say that one ought not to give to charities that we have a personal attachment to. It's just to say that, if one gives to charities that one has a personal attachment to that aren't the most effective ways of doing good, that doesn't discharge one's duties of beneficence.

4. How can we maximise our impact?

So far, I've mostly discussed health interventions. In those cases, we have a fairly well-established method for comparing effectiveness, namely years of healthy life (or QALYs) per dollar. But you might reasonably object that not all ways of helping others can be compared so easily. For example, the QALY framework doesn't allow us to evaluate interventions that aim to improve animal welfare, or ones that aim for political change, even though those could plausibly be among the most effective ways of doing good. This is a legitimate concern, and effective altruists have developed an alternative heuristic framework for prioritising among causes, including when the impact within some of those causes is difficult to measure. According to this framework, the following factors are indicative of which causes are highest-priority:²²

1. *Scale*. What's the magnitude of this problem? How much does it affect lives in the short run and long run?
2. *Neglectedness*. How many resources are already being dedicated to tackling this problem? How well allocated are the resources that are currently being dedicated to the problem? Is there reason to expect this problem can't be solved by markets or governments?
3. *Tractability*. How easy is it to make progress on this problem, and how easy is it to tell if you're making progress? Do interventions within this cause exist, and how strong is the evidence behind those interventions? Do you expect to be able to discover new promising interventions within this cause?

By way of illustration, let us briefly consider three cause areas that effective altruists have found promising.

Farm Animal Welfare

Every year 60 billion animals are killed for food,²³ and the vast majority of these spend their lives in factory farms in horrendous conditions; crammed together with little space, natural light or stimuli, and at constant risk of developing problems such as weakened or broken bones,

²² MacAskill, *Doing Good Better*

²³ https://www.ciwf.org.uk/media/3640540/ciwf_strategic_plan_20132017.pdf

infections and organ failure. Most have their lives ended prematurely when they are slaughtered for food. Farm animal welfare is highly *neglected*. Despite the size of the problem, in the US only a few tens of millions of philanthropic dollars are donated every year to organizations that focus on improving the lives of farmed animals—0.01% of total US philanthropy.²⁴ The amount is tiny even compared to other animal causes: there are 3000 times more animals in factory farms than there are stray pets, but factory farming gets one-fiftieth of the funding.²⁵ For this reason, additional funding can have a huge impact in this area, for example by enabling activists to campaign for large retailers and fast food chains to cut caged eggs out of their supply chains. According to research by the Open Philanthropy Project, one can spare somewhere between 38 and 250 hens a year of cage confinement per dollar spent.²⁶ Because of the sheer numbers of sentient beings involved, making progress on improving farm animal welfare could avert a huge amount of suffering.

Global health and development

Interventions within global health and development, such as alleviating extreme poverty and fighting infectious diseases, are considered a high priority within the effective altruism community, as they're particularly *tractable*. That is, efforts in global health have a great track record of improving lives.²⁷ Many interventions in these areas are incredibly cost-effective, too: as we've seen, on average it costs around \$3500 to save a life by distributing bednets.²⁸ Even simply transferring money to people who are very poor is a relatively cost-effective way of helping people.²⁹ Other top-recommended interventions include deworming schoolchildren, seasonal malaria chemoprevention, and vitamin A fortification.³⁰

Existential Risk

An *existential risk* is a risk of an event that would either annihilate intelligent life on Earth or permanently and drastically curtail its potential; such risks could come from all-out nuclear war, or extreme climate change, or a man-made global pandemic. Though low in probability, these are extremely great in *scale*: in addition to the deaths of all 7.5 billion people on this planet, an existential catastrophe would also entail the loss of all of humanity's future potential. If we avoid existential catastrophe, human civilisation could survive for hundreds of millions of years before the Earth is no longer habitable. The extinction of the human race would therefore mean the

²⁴ <http://s-risks.org/altruism-numbers-and-factory-farms/>

²⁵ <http://s-risks.org/altruism-numbers-and-factory-farms/>

²⁶ <https://www.openphilanthropy.org/blog/initial-grants-support-corporate-cage-free-reforms>

²⁷ <http://millionssaved.cgdev.org/>

²⁸ <https://www.givewell.org/how-we-work/our-criteria/cost-effectiveness/cost-effectiveness-models>

²⁹ "Economic Empowerment" *Giving What We Can*, April 2016.

<https://www.givingwhatwecan.org/report/cash-transfers/>

³⁰ <https://www.givewell.org/charities/top-charities>

preclusion of trillions of lives to come. Mitigating these risks may therefore be one of the most important moral issues we face.³¹

These three areas—animal welfare, global development, and existential risk—are currently among the top priorities for many effective altruists. However, as more evidence comes in, we are sure to discover new and important cause areas.

³¹ <https://80000hours.org/articles/extinction-risk/>