1

The Definition of Effective Altruism

William MacAskill

There are many problems in the world today. Over 750 million people live on less than $1.90 per day (at purchasing power parity).[1] Around 6 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] More than 3,000 nuclear warheads are in high-alert ready-to-launch status around the globe.[5] Bacteria are becoming antibiotic-resistant.[6] Partisanship is increasing, and democracy may be in decline.[7]

Given that the world has so many problems, and that these problems are so severe, surely we have a responsibility to do something about them. But what? There are countless problems that we could be addressing, and many different ways of addressing each of those problems. Moreover, 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. There are now thousands of people around the world who have chosen their careers, at least in part, on the basis of effective altruist ideas: individuals have gone into scientific research, think tanks, party politics, social entrepreneurship, finance (in order to do good through donating), and non-profit work.[8] Every year, over a thousand people in total gather at various Effective Altruism Global conferences, in locations as diverse as San Francisco, London, Hong Kong, and Nairobi.[9] Over 3,500 people have taken Giving What We Can's pledge to give at least 10 per cent of their income for the rest of their lives to

[1] World Bank Group (2016, ch. 2).   [2] UNICEF (2017).
[3] Broome (2012); Nordhaus (2015).
[4] United Nations Department of Economic and Social Affairs (2015).   [5] Davenport (2018).
[6] World Health Organization (2016).   [7] Norris and Inglehart (2018).
[8] For more information on effective altruism as applied to career choice, see www.80000hours.org.
[9] See www.eaglobal.org.

the organizations they believe to be most cost-effective, together pledging over $1.5 billion of lifetime donations.[10] Individuals donate over $90 million per year to GiveWell's top recommended charities,[11] and GoodVentures, a foundation that currently has potential assets of $14 billion, is committed to effective altruist principles and is distributing over $200 million each year in grants, advised by the Open Philanthropy Project.[12]

As a result of this, the effective altruism community has contributed to significant achievements in the areas of global catastrophic risk reduction, farm animal welfare, and global health. In 2016 alone, the effective altruism community was responsible for protecting 6.5 million children from malaria by providing long-lasting insecticide treated bednets, sparing 360 million hens from living in caged confinement, and providing significant impetus and support in the development of technical AI safety as a mainstream area of machine learning research.[13]

This movement has also inspired significant academic discussion. Books on the topic include *The Most Good You Can Do* by Peter Singer and my own *Doing Good Better*;[14] academic articles on effective altruism, both supportive and critical, have appeared in *Philosophy and Public Affairs*, *Utilitas*, *Journal of Applied Philosophy*, *Ethical Theory and Moral Practice*, and other publications.[15] A volume of *Essays in Philosophy* is dedicated to the topic and there is discussion of effective altruism by academics in the *Boston Review*.[16]

However, if we are to have a meaningful academic debate about effective altruism, we need to agree on what we're talking about. This chapter aims to help with that aim, introducing the Centre for Effective Altruism's definition, explaining why the Centre has chosen the definition it has, and providing a precise philosophical interpretation of that definition. I believe that this understanding of effective altruism, which is widely endorsed by those within the effective altruism community, is quite different from the understanding of effective altruism possessed by many in the general public and by many critics of effective altruism. In this essay, I explain why I prefer the definition I give, and then use the opportunity to correct some prevalent misunderstandings of effective altruism.

Before we begin, it's important to note that, in defining 'effective altruism', we are not attempting to describe a fundamental aspect of morality. In empirical research fields, we can distinguish between science and engineering. Science is the attempt to discover general truths about the world we live in. Engineering is

---

[10] 'The Giving What We Can Pledge.' Available at https://www.givingwhatwecan.org/pledge
[11] 'GiveWell's Impact.' Available at https://www.givewell.org/about/impact.
[12] 'How Can We Accomplish as Much Good as Possible?' Available at https://www.openphilanthropy.org/
[13] Bollard (2016); Dewy (2015).   [14] Singer (2015a); MacAskill (2015a).
[15] For example Berkey (2018); Pummer (2016); Gabriel (2017); MacAskill (2014); McMahan (2016).
[16] Singer (2015b).

the use of our scientific understanding to design and build structures or systems that benefit society.

We can make the same distinction within moral philosophy. Typically, moral philosophy is concerned with discovering general truths about the nature of morality—the equivalent of normative science. But there is also scope for the equivalent of engineering within moral philosophy, for example by creating new moral concepts whose use, if taken up broadly by society, would improve the world.

Defining 'effective altruism' is a matter of engineering rather than of describing some fundamental aspect of morality. In this vein, I suggest two principal desiderata for the definition. The first is to match the actual practice of those who are currently described as engaging in effective altruism, and the understanding of effective altruism that the leaders of the community have. The second is to ensure that the concept has as much public value as possible. This means, for example, we want the concept to be broad enough to be endorsable by or useful to many different moral views, but still determinate enough to enable users of the concept to do more to improve the world than they otherwise would have done. This, of course, is a tricky balancing act.

1. Previous definitions of effective altruism

The term 'effective altruism' was coined through the founding of the Centre for Effective Altruism, in a democratic process among seventeen people involved in the organization, on 3 December 2011.[17] However, no official definition of the term was introduced. Over the years, effective altruism has been defined in a number of distinct ways by different people. Here are some examples:

(1) To us, "effective altruism" means trying to do as much good as possible with each dollar and each hour that we have.[18]
(2) Effective altruism is about asking, "How can I make the biggest difference I can?" and using evidence and careful reasoning to try to find an answer.[19]
(3) Effective altruism is based on a very simple idea: we should do the most good we can…Living a minimally acceptable ethical life involves using a substantial part of our spare resources to make the world a better place. Living a fully ethical life involves doing the most good we can.[20]

---

[17] These people were: Will MacAskill (then 'Crouch'), Toby Ord, Nick Beckstead, Michelle Hutchinson, Holly Morgan, Mark Lee, Tom Ash, Matt Wage, Ben Todd, Tom Rowlands, Niel Bowerman, Robbie Shade, Matt Gibb, Richard Batty, Sally Murray, Rob Gledhill, and Andreas Mogensen.

[18] Karnofsky (2013).   [19] MacAskill (2015a, pp. 14–15).   [20] Singer (2015b).

(4) Effective altruism is a research field which uses high-quality evidence and careful reasoning to work out how to help others as much as possible. It is also a community of people taking these answers seriously, focusing their efforts on the most promising solutions to the world's most pressing problems.[21]

(5) Effective altruism is a philosophy and social movement that uses evidence and reason to determine the most effective ways to benefit others.[22]

We can see some commonalities among these definitions.[23] All invoke the idea of maximization, and all are about the achievement of some value, whether that's the value of increasing wellbeing, or simply of achieving the good in general. However, there are differences, too. Definitions (1)–(3) talk about 'doing good' whereas definitions (4) and (5) talk about 'helping others' and 'benefitting others'. Unlike the others, definition (3) makes effective altruism a normative claim, rather than a non-normative project, such as an activity or research field or movement. Definitions (2), (4), and (5) invoke the idea of using evidence and careful reasoning, whereas definitions (1) and (3) do not.

The Centre for Effective Altruism's definition takes a stand on each of these issues, defining effective altruism as follows: effective altruism is about using evidence and reason to figure out how to benefit others as much as possible, and taking action on that basis.[24]

I led on the creation of this definition, with input from a wide number of advisors in the effective altruism community, and significant help from Julia Wise and Rob Bensinger. It and a set of guiding values that sit alongside it have been formally endorsed by the large majority of leaders in the effective altruism community.[25] There is no 'official' definition of effective altruism, but the Centre's definition is closer to being one than any other. However, this statement of effective altruism was intended for a general rather than a philosophical audience, so some

---

[21] 'Introduction to Effective Altruism' (2016).   [22] 'Effective Altruism.' *Wikipedia*.

[23] I'll treat each of these as definitions, although only the fourth had the right grammatical form to be one. All these statements are intended to be read by a general audience, so I don't place much weight on specific word choice like 'is about' or 'is based on'.

[24] This definition is accompanied by a set of guiding principles, that are intended to form a broad code of conduct for those in the effective altruism community. These principles are: commitment to others, scientific mindset, openness, integrity, and collaborative spirit. See 'CEA's Guiding Principles', Centre for Effective Altruism, https://www.centreforeffectivealtruism.org/ceas-guiding-principles/.

[25] This includes the following organizations: Impact Investing, 80,000 Hours, Animal Charity Evaluators, Charity Science, Effective Altruism Foundation, Foundational Research Institute, Future of Life Institute, Raising for Effective Giving, and The Life You Can Save. And it includes the following individuals (though not their respective organisations): Elie Hassenfeld of GiveWell, Holden Karnofsky of the Open Philanthropy Project, Toby Ord of the Future of Humanity Institute, Peter Singer of Princeton University and the University of Melbourne, and Nate Soares of the Machine Intelligence Research Institute.

precision was lost for the sake of accessibility. For that reason, I'd like to provide and then unpack a more precise formulation here. My definition is as follows:

Effective altruism is:

(*i*)   the use of evidence and careful reasoning to work out how to maximize the good with a given unit of resources, tentatively understanding 'the good' in impartial welfarist terms, and

(*ii*)   the use of the findings from (*i*) to try to improve the world.

(*i*)   refers to effective altruism as an intellectual project (or 'research field'); (*ii*) refers to effective altruism as a practical project (or 'social movement').
  The definition is:

- *Non-normative.* Effective altruism consists of two projects, rather than a set of normative claims.
- *Maximizing.* The point of these projects is to do as much good as possible with the resources that are dedicated towards it.
- *Science-aligned.* The best means to figuring out how to do the most good is the scientific method, broadly construed to include reliance on careful rigorous argument and theoretical models as well as data.
- *Tentatively impartial and welfarist.* As a tentative hypothesis or a first approximation, doing good is about promoting wellbeing, with everyone's wellbeing counting equally. More precisely: for any two worlds A and B with all and only the same individuals, of finite number, if there is a one-to-one mapping of individuals from A to B such that every individual in A has the same wellbeing as their counterpart in B, then A and B are equally good.[26]

I'll explain why these choices were made, in turn.
  Two of the choices are uncontroversial. First, every proposed definition of effective altruism is maximizing, and this idea is baked into almost every explanation of effective altruist ideas, including the title of Peter Singer's book *The Most Good You Can Do*. However, there is an ambiguity that needs to be clarified. One can try to increase the amount of good one does in two ways: by increasing the amount of resources that one dedicates to doing good; and by trying to increase the effectiveness of the resources that one has dedicated to doing good. On the definition I suggest, effective altruism is about maximizing only in the

---

[26] Note that, read literally, the use of "benefit others" in CEA's definition would rule out some welfarist views, such as the view on which one can do good by creating good lives but that this does not involve benefiting those who would otherwise not exist. In this case, philosophical precision was sacrificed for readability.

latter sense. On other definitions this has not been clear; I explain the reasons for this choice in the next section.

Second, the idea that effective altruism involves relying on the scientific method, broadly construed, is also clearly a core part of the concept. All the major research organizations within effective altruism involve relying on data or scientific research where it is possible to do so, as well as on theoretical models and on clear and rigorous argument.

Again, however, a clarification is warranted. Sometimes critics interpret effective altruism's endorsement of the scientific method to mean that we rely solely on randomized controlled trials (RCTs). This, if true, would of course be naïve. But we should understand the 'scientific method' much more broadly than that. There are some issues that, for practical reasons, we cannot assess directly on the basis of an RCT, such as what the probability of human extinction is over the next two centuries. There are also a wide variety of ways of gaining empirical evidence other than RCTs, such as regressions, quasi-experiments, surveys, and simple fact-finding. And there are many issues for which experimental evidence in general is not relevant, such as in ethics, epistemology, and decision theory.

The two more controversial aspects of the definition are that it is non-normative, and that it is tentatively impartial and welfarist. I'll discuss these in turn.

2. Effective altruism as a project, rather than a normative claim

The definition of effective altruism I've given presents effective altruism as consisting of two projects: an intellectual project, of trying to figure out how to use resources in whatever way will do the most good with a given unit of resources; and a practical project, of putting the results from the intellectual project into practice and trying to use some of one's resources to improve the world.

There are two ways in which the definition of effective altruism could have made normative claims. First, it could have made claims about how much one is required to sacrifice: for example, it could have stated that everyone is required to use as much of their resources as possible in whatever way will do the most good; or it could have stated some more limited obligation to sacrifice, such as that everyone is required to use at least 10 per cent of their time or money in whatever way will do the most good.

There were three reasons why we didn't include an obligation to sacrifice in the definition. First, it was very unpopular among leaders of the effective altruism community: in a survey of such leaders in 2015, 80 per cent of respondents stated that they thought the definition should not include a sacrifice component and only 12.5 per cent thought it should contain a sacrifice component. Second, within the effective altruism community more broadly, only some members believe that

one has an obligation to engage in effective altruism; others believe that engaging in effective altruism is part of a meaningful life for them, but that there is no obligation to do so. A 2017 survey of 1,843 members of the effective altruism community included the question, 'Do you think of Effective Altruism more as an "opportunity" or an "obligation"?' In response, 56.5 per cent chose 'moral duty' or 'obligation', and 37.7 per cent chose 'opportunity' (there was no option in that year to choose 'both').[27] In the previous effective altruism survey, in 2015, 42 per cent of respondents chose 'both' in response to the same question, 34 per cent chose 'opportunity' and 21 per cent chose 'obligation'.[28]

Third, it makes the concept far more ecumenical. Because effective altruism is not a normative claim, it's consistent with any moral view. But the project is still of interest to those with many different moral views: most plausible moral views would allow that there is a *pro tanto* reason to promote the good, and that wellbeing is of some value, and therefore that the question of how one can do the most to promote welfarist value with a given unit of resources needs to be resolved as one part of answering the question of how to live a morally good life. In contrast, any sort of claim about our obligations to maximize the good will be more controversial, particularly if we try to make a general statement covering people of very different income levels and personal situations. The public value of the concept of effective altruism therefore seems greater if it does not include a sacrifice component: it allows a wider range of people to engage in effective altruism, preventing the concept from being off-putting to those who don't believe that there are strong obligations of beneficence, in general or in their particular case. This is backed up by the anecdotal experience of those involved in Giving What We Can: those in the organization initially tried out both 'obligation' and 'opportunity' framings to encourage people to take the 10 per cent pledge, finding that the 'opportunity' framing was much more efficacious. This fact could also explain why Giving What We Can caused such a rise in the number of people taking Peter Singer's views on our obligations of beneficence very seriously, despite these ideas being around for decades prior.

Finally, it focuses attention on the most distinctive aspect of effective altruism: the open question of how we can use resources to improve the world as much as possible. This question is much more neglected and arguably more important than the question of how much and in what form altruism is required of one.[29] For this reason, most people within the effective altruism community are much more concerned with getting on with the project of figuring out *how* we can do

---

[27] McGeoch and Hurford (2017). Note that the sample was non-random: everyone who wanted to answer the survey was able to, and it was advertised as widely as possible within the community. As a result, all statistics drawn from this survey should be taken as suggestive but not definitive.

[28] Cundy (2015).

[29] For an argument that it is more important, see Chapter 2 of this volume by Toby Ord.

more good rather than asking to what extent, or in what way, we are required to do the most good.

The second way in which we could have made the definition normative is by appeal to conditional obligations. For example, the definition could have included the idea that if one is trying to use resources to do good, one ought to choose whatever action will maximize the good, subject to not violating any side constraints.[30]

I think that the case for being non-normative in this sense is not as strong as the case against including a sacrifice component, but we kept the definition entirely non-normative for much the same reasons that we did not want to include a sacrifice component. First, most EA leaders were against it: in the 2015 survey, 70 per cent of respondents stated that they thought the definition should be non-normative and only 20 per cent thought it should be normative.

Second, again, is ecumenicism. There are reasonable views on which, because it's permissible whether to use one's resources to do good, it's also permissible to aim to do some good, but less good than one could have done. Moreover, even if we think that sometimes conditional obligations of this form hold, there are also difficult questions about the scope of such obligations. We clearly would not want to commit to there being a conditional obligation to maximize the good in cases where doing so would violate someone's rights, but what about in conditions where it would violate the actor's integrity? Or in cases where one has already spent most of one's resources altruistically, but now wants to spend some of one's money on charities that are less effective but dear to one's heart? Any view on this topic will be highly controversial.[31]

We could dilute the normative claim by phrasing it merely in terms of reasons, for example, stating merely that one has some reason to do as much good as possible. But if so, then effective altruism would be a very weak claim, and not a very interesting one. The distinctive aspect of effective altruism is the choice to focus on asking how we can use some of our resources to do as much good as possible, and the conclusions we come to about how to do as much good as possible, not the very thin claim that one has some reason to do as much good as possible.

3. Effective altruism as tentatively impartial and welfarist

The second controversial part of the definition is that it is tentatively impartial and welfarist. It is tricky to delineate which axiological views should be counted as within the remit of effective altruism, and which should be counted as outside

---

[30] The idea of conditional obligations is explored by Pummer (2016), though the claim he defends is significantly weaker than this.

[31] See, for example, Pummer (2016); Sinclair (2018); McMahan (2018).

of effective altruism. On one end of the spectrum, we could define effective altruism as the attempt to do the most good, according to whatever view of the good the individual in question adheres to. On the other end of the spectrum, we could define effective altruism as the attempt to do the most good on one very particular understanding of the good, such as total hedonistic utilitarianism. Either choice faces severe problems. If we allow any view of the good to count, then white supremacists could count as practicing effective altruism, which is a conclusion that we clearly do not want. If we restrict ourselves to one particular view of the good, then we lose any claim to ecumenicism, and we also misrepresent the effective altruism community itself, which has vibrant disagreement over many areas of axiology.

Alternatively, one could attempt to restrict effective altruism to cover only 'reasonable' views of the good. But then, first, we face the difficulty of explaining what counts as 'reasonable'. And, second, we also misrepresent the practices of the effective altruism community, which is distinctive insofar as it is currently so focused on wellbeing, and insofar as all the analyses from the leading effective altruist research organizations count each individual's interests equally. What's more, I think that it is unlikely in the foreseeable future that the community will have people or projects focusing, for example, on art or biodiversity as ends in themselves. Similarly, it is unlikely that those in the community would focus on rectifying injustice in cases where they believed that there were other available actions which, though they would leave the injustice remaining, would do more good overall.

My preferred solution is tentative impartial welfarism, defined above. This excludes non-welfarist views on which, for example, biodiversity or art has intrinsic value, and excludes partialist views on which, for example, the wellbeing of one's co-nationals count for more than those of foreigners. But it includes utilitarianism, prioritarianism, sufficientarianism, egalitarianism, different views of population ethics, and different views of how to weight the wellbeing of different creatures.

This welfarism is 'tentative', however, insofar as it is taken to be merely a working assumption. The ultimate aim of the effective altruist project is to do as much good as possible; the current focus on wellbeing rests on the idea that, given the current state of the world and our incredible opportunity to benefit others, the best ways of promoting welfarist value are broadly the same as the best ways of promoting the good. If that view changed and those in the effective altruism community were convinced that the best way to do good might well involve promoting non-welfarist goods, then we would revise the definition to simply talk about 'doing good' rather than 'benefiting others'.

I believe that this understanding is supported by the views of EA leaders. In the 2015 survey of EA leaders referred to earlier, 52.5 per cent of respondents were

in favour of the definition including welfarism and impartiality, with 25 per cent against. So the inclusion of impartial welfarism has broad support, but not as convincing support as other aspects of the definition.

What's more, this restriction does little to reduce effective altruism's ecumenicism: wellbeing is part of the good on most or all plausible moral views. Effective altruism is not claiming to be a complete account of the moral life. But, for any view that takes us to have reasons to promote the good, and that says wellbeing is part of the good, the project of working out how we can best promote wellbeing will be important and relevant.

Having explained what effective altruism *is*, let's now turn to what effective altruism *is not*, and address some common misconceptions.

4. Misunderstandings of effective altruism

4.1 Misconception #1: Effective altruism is just utilitarianism

Effective altruism is often considered to simply be a rebranding of utilitarianism, or to merely refer to applied utilitarianism. John Gray, for example, refers to 'utilitarian effective altruists', and in his critique does not distinguish between effective altruism and utilitarianism.[32] Giles Fraser claims that the 'big idea' of effective altruism is 'to encourage a broadly utilitarian/rationalist approach to doing good.'[33]

It is true that effective altruism has some similarities with utilitarianism: it is maximizing, it is primarily focused on improving wellbeing, many members of the community make significant sacrifices in order to do more good, and many members of the community self-describe as utilitarians.[34]

But this is very different from effective altruism being the same as utilitarianism. Unlike utilitarianism, effective altruism does not claim that one must always sacrifice one's own interests if one can benefit others to a greater extent.[35] Indeed, on the above definition effective altruism makes *no* claims about what obligations of benevolence one has.

---

[32] Gray (2015).   [33] Fraser (2017); Bakić (2015); Gabriel (2015); Tumber (2015).

[34] In the 2017 effective altruism survey, 52.8 per cent of respondents chose 'utilitarianism' in response to the question 'What moral philosophy, if any, do you lean towards?' In addition, 12.6 per cent chose 'consequentialism (NOT utilitarianism)', 5.2 per cent chose 'virtue ethics', 3.9 per cent chose 'deontology', and 25.5 per cent chose 'no opinion, or not familiar with these terms'. As a caveat, however, it's not clear how well the respondents understood these terms. For example, in conversation I learned that one respondent thought that utilitarianism refers to any moral theory that can be represented by a utility function.

[35] On the demandingness objection to utilitarianism, see 'The Demandingness of Morality: Toward a Reflective Equilibrium' (Berkey 2016).

Unlike utilitarianism, effective altruism does not claim that one ought always to do the good, no matter what the means;[36] indeed, as suggested in the guiding principles, there is a strong community norm against 'ends justify the means' reasoning. This is emphasized, for example, in an 80,000 Hours blog post by Ben Todd and I.[37]

Finally, unlike utilitarianism, effective altruism does not claim that the good equals the sum total of wellbeing. As noted above, it is compatible with egalitarianism, prioritarianism, and, because it does not claim that wellbeing is the only thing of value, with views on which non-welfarist goods are of value.[38]

In general, very many plausible moral views entail that there is a *pro tanto* reason to promote the good, and that improving wellbeing is of moral value.[39] If a moral view endorses those two ideas, then effective altruism is part of the morally good life.

4.2 Misconception #2: Effective altruism is just about fighting poverty

The vast majority of the focus on effective altruism in the media and in critical academic discussion has been on the part of effective altruism that is about fighting poverty. For example, Judith Lichtenberg begins her article with the question, "How much money, time, and effort should you be giving to relieve dire poverty?"[40] Jennifer Rubenstein describes effective altruism as "a social movement focused on alleviating poverty," and Iason Gabriel describes effective altruism as encouraging "individuals to do as much good as possible, typically by contributing money to the best-performing aid and development organisations."[41]

It is, of course, true that fighting poverty is one core focus of those in the effective altruism community. In the 2017 EA survey, 41 per cent of respondents identified extreme poverty as their top priority cause area, and some effective altruist organizations such as GiveWell are exclusively focused on poverty alleviation[42] (just as some other organizations within effective altruism are focused exclusively on animal welfare[43] or existential risks).[44]

But two core parts of effective altruism are *cause-neutrality* and *means-neutrality*: being open in principle to focusing on any problem (such as global health, or climate change, or factory farming) and being open in principle to using any (non-side-constraint violating) means to addressing that problem. In every

[36] On utilitarianism and constraints, see Kagan (1989).
[37] Todd and MacAskill (2017).   [38] See Parfit (1997); Temkin (1993); Hurka (1993).
[39] Kagan (1998); Ross (1930).   [40] Lichtenberg (2015).   [41] Gabriel (2017).
[42] McGeoch and Hurford (2017).   [43] For example, Animal Charity Evaluators (ACE).
[44] For example, the Berkeley Existential Risk Initiative (BERI).

case, the criterion is simply what activity will do the most good. Cause and means neutrality follow straightforwardly from the assumptions of maximization and impartial welfarism. If, by focusing on one cause rather than another, or by choosing one means rather than another, one can do more to promote wellbeing (without violating any side constraints) then someone who is committed to effective altruism will do so.

And, in practice, members of the effective altruism community support many other causes, including animal suffering reduction, criminal justice reform, and existential risk mitigation. In the 2017 EA survey, in addition to the 41 per cent of respondents who identified extreme poverty as their top priority cause area, 19 per cent of respondents chose cause prioritization as the top priority, 16 per cent chose AI, 14 per cent chose environmentalism, 12 per cent chose promoting rationality, 10 per cent chose non-AI existential risk, and 10 per cent chose animal welfare. These results were broadly similar to the 2015 and 2014 surveys: poverty is the most common focus area for individuals in the effective altruism community, but is not the focus for the majority of individuals in the community.

This is mirrored when we look at the distribution of grants by the Open Philanthropy Project. In 2017, they spent:

- $118 million (42 per cent) on global health and development
- $43 million (15 per cent) on potential risks from advanced artificial intelligence
- $36 million (13 per cent) on scientific research (which cuts across other causes)
- $28 million (10 per cent) on biosecurity and pandemic preparedness
- $27 million (10 per cent) on farm animal welfare
- $10 million (4 per cent) on criminal justice reform
- $9 million (3 per cent) on other global catastrophic risks
- $10 million (4 per cent) on other cause areas, including land use reform, macroeconomic policy, immigration policy, promotion of effective altruism, and improving decision-making

The amount of money received by the Effective Altruism Funds—where individual donors can give to a fund managed by an expert for regranting within a particular cause area—tells a similar story. In 2017 it received:

- $982,000 (48 per cent) for the global health and development fund
- $409,000 (20 per cent) for the animal welfare fund
- $363,000 (18 per cent) for the long-term future fund
- $290,000 (14 per cent) for the effective altruism community fund

So, in contrast to the equation of effective altruism with poverty reduction only, a more accurate description would be that the effective altruism community

currently focuses on extreme poverty, factory farming, and existential risk, with a small number of other areas of focus.

4.3 Misconception #3: Effective altruism is entirely about donations or earning to give

Most media attention focuses on the part of effective altruism that focuses on effective altruism as applied to donations, and a significant proportion has focused on the idea of 'earning to give'—that people should deliberately pursue a lucrative career in order to be able to donate a large proportion of those earnings to effective charities.[45]

This is also true for the criticism of effective altruism. Iason Gabriel described effective altruism as 'a philosophy and social movement that aims to revolutionise the way we do philanthropy', and focuses his discussion on effective altruism and charity.[46] Similarly, Jennifer Rubenstein's review of *Doing Good Better* and *The Most Good You Can Do* focuses on the charitable side of the effective altruism movement.[47]

There's no doubt that philanthropy is a major focus of the effective altruism community, and 80,000 Hours recognize that they promoted earning to give too heavily in their early marketing materials,[48] and so it's entirely reasonable for an article to focus on that aspect. But it means that a casual observer could think that this is *all* that the effective altruism focuses on, even though it is not the only focus.

The organization 80,000 Hours is entirely focused on helping individuals to use their career as effectively as possible. And they recommend that only about 15 per cent of altruistic graduates who would be happy in a wide variety of career paths should earn to give in the long term.[49] Similarly, in large part because of the success of the EA movement at raising philanthropic money, the primary focus of the Centre for Effective Altruism is to encourage people to move into working in particularly important causes, rather than funding those causes.[50] And in the 2015 EA survey, survey-takers were asked, 'What broad career path are you planning to follow?' Although earning to give was the most common response, receiving 36 per cent of responses, 13 per cent selected 'non-profit' work, 25 per cent selected 'research', and 26 per cent selected 'none of these'. It seems that most members of the effective altruism community, therefore, do not plan to use donations as their main path to impact.

---

[45] For examples, see Herzog (2016); Rubenstein (2015); Earle and Read (2016); and my own article arguing in favour of this position is 'Replaceability, Career Choice, and Making a Difference' (2014).
[46] Gabriel (2017).   [47] Rubenstein (2015).   [48] 'Our Mistakes' 80,000 Hours.
[49] MacAskill (2015b).   [50] Hesketh-Rowe (2017).

4.4 Misconception #4: Effective altruism
ignores systemic change

Of all the criticisms of effective altruism, the most common is that effective altruism ignores systemic change. For example, Brian Leiter comments that: "I am a bit skeptical of undertakings like [effective altruism], for the simple reason that most human misery has systemic causes, which charity never addresses, but which political change can address; ergo, all money and effort should go towards systemic and political reform."[51] This objection is also discussed by Amia Srinivasan,[52] Iason Gabriel,[53] and Jennifer Rubenstein.[54]

But effective altruism is clearly open to systemic change in both principle and practice.[55] We can distinguish a broader and a narrower sense of 'systemic change'. In the broader sense, a systemic change is any change that involves a one-off investment in order to reap a long-lasting benefit. In the narrower sense, 'systemic change' refers to long-lasting *political* change. Either way, the allegation is often that those in the effective altruism community have been biased by a desire for quantification away from difficult-to-assess measures such as political change.[56]

It's clear that effective altruism is open to systemic change in principle: effective altruism is committed to cause-neutrality and means-neutrality, so if improving the world in some systemic way is the course of action that will do the most good (in expectation, without violating any side constraints), then it's the best course of action by effective altruism's lights. More importantly, however, effective altruists often advocate for systemic change in practice, even in the narrower sense. An incomplete list of examples is as follows:[57]

- International labour mobility has been a focus area of members of the effective altruism community for some time. Openborders.info, run by a member of the effective altruism community, collates research on and promotes the option of dramatic increases in migration from poor to rich countries. Open Philanthropy has made grants in this area, including to the Center for Global Development, the US Association for International Migration, and ImmigrationWorks. The reason for this focus is that one of the structural reasons why people in poor countries are poor is that they are unable to move to countries where they could be more productive. In effect, they are being incarcerated in the country into which they were born by the joint migration restrictions of all other countries. For this reason, there are

---

[51] Leiter (2015).   [52] Srinivasan (2015).   [53] Gabriel (2017).
[54] Rubenstein (2015). Other instances of this criticism include Herzog (2016); Snow (2015); Earl and Read (2016). [See also Gabriel and McElwee, Chapter 7 of this volume].
[55] For further discussion of this issue, see Berkey (2018).   [56] Clough (2015).
[57] For further discussion, see Wiblin (2015).

economic arguments that the benefits to people in poverty from greater freedom of movement across borders would be enormous.[58]

- The Center for Election Science promotes alternative voting systems, in particular approval voting; it's run by a member of the effective altruism community, and received a grant from the Open Philanthropy Project at my recommendation.[59]

- The Centre for Effective Altruism has provided advice for the World Bank, the WHO, the Department for International Development, and Number 10 Downing Street.

- 80,000 Hours' list of recommended careers includes party politics, policy-oriented civil service, and think tanks, and has an employee entirely dedicated to advising people who wish to work in policy and government in the area of technological risk.

- The animal welfare wing of the effective altruism community, including Mercy for Animals and The Humane League, has had astonishing success by lobbying large retailers and fast food chains to get them to pledge to no longer use eggs from caged hens in their supply chain.

- Organizations such as the Future of Humanity Institute and the Centre for the Study of Existential Risk are actively working on policy around developments of new technology, and advising organizations such as the US government, UK government and the UN.

- The Open Philanthropy Project has made numerous grants within the areas of land use reform, criminal justice reform, improving political decision-making, and macroeconomic policy.[60]

Once we consider the broader sense of systemic change, then an even larger proportion of effort from the effective altruism community is focused on systemic change. For example, all work addressing existential risks is in this category, as is the focus on scientific research and on improving science (such as through encouraging preregistration of trials), as is the focus on developing lab-grown meat and plant-based meat substitutes.

Of course, it's perfectly plausible that there are 'systemic' interventions that those in the effective altruism community are neglecting. Perhaps campaigning to create an international law banning the purchase of natural resources from dictatorships is an even more effective activity than any of the current activities of effective altruists.[61] But this is an in-house dispute, rather than a criticism of effective altruism per se. One could argue that it's in the nature of the way of

---

[58] Caplan and Naik (2015, ch. 8).
[59] See this introduction to voting theory by a board member of the Center for Election Science: Quinn (2018).
[60] Grant Database. Open Philanthropy Project.
[61] See Chapter 7 of this volume, 'Effective Altruism, Global Poverty, and Systemic Change'.

thinking of those in the effective altruism community that this idea is neglected. But there are ready alternative explanations: the chance of such a campaign being successful is astronomically low and, even if it were successful, even in the best case scenarios the legal change would occur decades hence, when the problem of extreme poverty will probably be far smaller and less severe than it is today.[62] Given this, and given the commitments to systemic change listed above, it's hard to see why we should think of this as a criticism of effective altruism per se, rather than simply a disagreement about the best ways of promoting wellbeing.

5. Conclusion

In this chapter, I've unpacked the Centre for Effective Altruism's definition of effective altruism, and explained some of the reasons why we chose that definition. I've then responded to some common misunderstandings of effective altruism. In doing so, I hope that I have helped to add clarity to future debates around effective altruism, allowing us to see which objections, if successful, would show that effective altruism has little or no place in our moral lives, and which are really just in-house debates about how to do the most good.

References

Bakić, Marko. 2015. 'How Is Effective Altruism Related to Utilitarianism?'Quora, December30.Availableathttps://www.quora.com/How-is-effective-altruism-related-to-utilitarianism.

Berkey, Brian. 2016. 'On the Demandingness Objection to Utilitarianism.' *Philosophical Studies* 173 (11): 3015–35.

Berkey, Brian. 2018. 'The Institutional Critique of Effective Altruism.' *Utilitas* 30 (2): 143–71.

Bollard, Lewis. 2016. 'Initial Grants to Support Corporate Cage-Free Reforms.' *Open Philanthropy Project* (blog), 31 March. Available at https://www.openphilanthropy.org/blog/initial-grants-support-corporate-cage-free-reforms.

Broome, John. 2012. *Climate Matters: Ethics in a Warming World*. New York: W.W. Norton.

Caplan, Bryan, and Vipul Naik. 2015. 'A Radical Case for Open Borders.' In *The Economics of Immigration: Market-Based Approaches, Social Science, and Public Policy*, Benjamin Powell, ed. New York: Oxford University Press, ch. 8.

---

[62] Poverty has decreased dramatically over the past two centuries, and we should expect this trend to continue. See Roser and Ortiz-Ospina (2017).

Clough, Emily. 2015. 'Effective Altruism's Political Blind Spot.' *Boston Review*, 14 July. Available at https://bostonreview.net/world/emily-clough-effective-altruism-ngos.

Cundy, Chris. 2015. 'The 2015 Survey of Effective Altruists: Results and Analysis.' Effective Altruism Forum, 29 July. Available at http://effective-altruism.com/ea/zw/the_2015_survey_of_effective_altruists_results/.

Davenport, Kelsey. 2018. 'Nuclear Weapons: Who Has What at a Glance.' Arms Control Association. Updated June 2018. Available at https://www.armscontrol.org/factsheets/Nuclearweaponswhohaswhat.

Dewey, Daniel. 2015. 'Potential Risks from Advanced Artificial Intelligence.' Open Philanthropy Project. Available at https://www.openphilanthropy.org/focus/global-catastrophic-risks/potential-risks-advanced-artificial-intelligence.

Earle, Sam, and Rupert Read. 2016. 'Why "Effective Altruism" is Ineffective: The Case of Refugees.' *Ecologist*, 5 April. Available at https://theecologist.org/2016/apr/05/why-effective-altruism-ineffective-case-refugees.

'Effective Altruism.' *Wikipedia* (Accessed 25 May 2018). Available at https://en.wikipedia.org/wiki/Effective_altruism.

Fraser, Giles. 2017. 'It's Called Effective Altruism—But is it Really the Best Way to Do Good?' *The Guardian*, 23 November. Available at https://www.theguardian.com/money/belief/2017/nov/23/its-called-effective-altruism-but-is-it-really-the-best-way-to-do-good.

Gabriel, Iason. 2015. 'The Logic of Effective Altruism.' *Boston Review* (forum review), 6 July. Available at https://bostonreview.net/forum/logic-effective-altruism/iason-gabriel-response-effective-altruism.

Gabriel, Iason. 2017. 'Effective Altruism and Its Critics.' *Journal Applied Philosophy* 34 (4): 457–73.

'GiveWell's Impact.' GiveWell (Accessed March 2018). Available at https://www.givewell.org/about/impact.

'The Giving What We Can Pledge.' Giving What We Can. Available at https://www.givingwhatwecan.org/pledge.

Gray, John. 2015. 'How & How Not to Be Good.' *The New York Review of Books*, 21 May. Available at http://www.nybooks.com/articles/2015/05/21/how-and-how-not-to-be-good/.

Herzog, Lisa. 2016. 'Can "Effective Altruism" Really Change the World?' openDemocracy, 22 February. Available at https://www.opendemocracy.net/transformation/lisa-herzog/can-effective-altruism-really-change-world.

Hesketh-Rowe, Larissa. 2017. 'CEA's 2017 Review and 2018 Plans.' Centre for Effective Altruism (blog), 18 December. Available at https://www.centreforeffectivealtruism.org/blog/cea-s-2017-review-and-2018-plans/.

'How Can We Accomplish as Much Good as Possible?' Open Philanthropy Project. Available at https://www.openphilanthropy.org/.

Hurka, Thomas. 1993. *Perfectionism*. New York: Oxford University Press.

Iason, Gabriel. 2017. 'Effective Altruism and Its Critics.' *Journal of Applied Philosophy* 34 (4): 457–73.

'Introduction to Effective Altruism.' 2016. Effective Altruism, 22 June. Available at https://www.effectivealtruism.org/articles/introduction-to-effective-altruism./

Kagan, Shelly. 1989. *The Limits of Morality*. New York: Oxford University Press.

Kagan, Shelly. 1998. *Normative Ethics*. Boulder, Colorado: Westview Press.

Karnofsky, Holden. 2013. 'Effective Altruism.' The GiveWell Blog, 13 August. Available at https://blog.givewell.org/2013/08/13/effective-altruism/.

Leiter, Brian. 2015. 'Effective Altruist Philosophers.' *Leiter Reports*, 22 June. Available at http://leiterreports.typepad.com/blog/2015/06/effective-altruist-philosophers.html.

Lichtenberg, Judith. 2015. 'Peter Singer's Extremely Altruistic Heirs.' *The New Republic*, 30 November. Available at https://newrepublic.com/article/124690/peter-singers-extremely-altruistic-heirs.

MacAskill, William. 2014. 'Replaceability, Career Choice, and Making a Difference.' *Ethical Theory and Moral Practice* 17 (2): 269–83.

MacAskill, William. 2015a. *Doing Good Better: How Effective Altruism Can Help You Make a Difference*. New York: Penguin Random House.

MacAskill, William. 2015b. '80,000 Hours Thinks that Only a Small Proportion of People Should Earn to Give Long Term.' 80,000 Hours, 6 July.

McGeoch, Ellen, and Peter Hurford. 2017. 'EA Survey 2017 Series: Distribution and Analysis Methodology.' Effective Altruism Forum, 29 August. Available at http://effective-altruism.com/ea/1e0/effective_altruism_survey_2017_distribution_and/.

McMahan, Jeff. 2016. 'Philosophical Critiques of Effective Altruism.' *The Philosophers' Magazine* 73: 92–9.

McMahan, Jeff. 2018. 'Doing Good and Doing the Best.' In *The Ethics of Philanthropy: Philosophers' Perspectives on Philanthropy*, Paul Woodruff, ed. New York: Oxford University Press, pp. 78–102.

Nordhaus, William. 2015. *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World*. New Haven: Yale University Press.

Norris, Pippa, and Ronald Inglehart. 2018. *Cultural Backlash: The Rise of Populist Authoritarianism*. New York: Cambridge University Press.

Parfit, Derek. 1997. 'Equality and Priority.' *Ratio* 10: 202–21.

Pummer, Theron. 2016. 'Whether and Where to Give.' *Philosophy & Public Affairs* 44 (1): 77–95.

Quinn, Jameson. 2018. 'A Voting Theory Primer for Rationalists.' LessWrong (blog), 12 April. Available at https://www.lesswrong.com/posts/D6trAzh6DApKPhbv4/a-voting-theory-primer-for-rationalists.

Ross, David. 1930. *The Right and the Good*. Oxford: Oxford University Press, ch. 2.

Roser, Max, and Ortiz-Ospina. 2017. 'Global Extreme Poverty.' Our World in Data, 27 March. Available at https://ourworldindata.org/extreme-poverty.

Rubenstein, Jennifer. 2015. 'The Logic of Effective Altruism.' *Boston Review* (forum review), 1 July. Available at: https://bostonreview.net/forum/logic-effective-altruism/jennifer-rubenstein-response-effective-altruism.

Sinclair, T. 2018. 'Are We Conditionally Obligated to Be Effective Altruists?' *Philosophy and Public Affairs* 46 (1): 36–59.

Singer, Peter. 2015a. *The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically*. New Haven, CT: Yale University Press.

Singer, Peter. 2015b. 'The Logic of Effective Altruism' and replies. *Boston Review*, 1 July. Available at https://bostonreview.net/forum/peter-singer-logic-effective-altruism.

Snow, Matthew. 2015. 'Against Charity.' *Jacobin*, 25 August.

Srinivasan, Amia. 2015. 'Stop the Robot Apocalypse.' *London Review of Books* 37 (18): 3–6.

Temkin, Larry. 1993. *Inequality*. New York: Oxford University Press.

Todd, Benjamin, and William MacAskill. 2017. 'Is it Ever Okay to Take a Harmful Job in Order to Do More Good? An In-depth Analysis.' 80,000 Hours. Available at https://80000hours.org/articles/harmful-career/.

Tumber, Catherin. 2015. 'The Logic of Effective Altruism.' *Boston Review*, 1 July. Availableathttps://bostonreview.net/forum/logic-effective-altruism/catherine-tumber-response-effective-altruism.

UNICEF. 2017. 'Levels & Trends in Child Mortality.' Available at https://www.unicef.org/publications/index_101071.html.

United Nations Department of Economic and Social Affairs. 2015. *The World's Women 2015: Trends and Statistics*. New York: United Nations.

Wiblin, Robert. 2015. 'Effective Altruists Love Systemic Change.' 80,000 Hours, 18 July. Available at https://80000hours.org/2015/07/effective-altruists-love-systemic-change/.

World Bank Group. 2016. *Poverty and Shared Prosperity: Taking on Inequality 2016*. Herndon: World Bank Publications.

World Health Organization. 2016. 'Antibiotic Resistance Key Facts.' Available at http://www.who.int/mediacentre/factsheets/antibiotic-resistance/en/.