

The Philosophical Core of Effective Altruism

Brian Berkey

1. Introduction

In his recent book on the effective altruism movement, Peter Singer approvingly cites a definition according to which effective altruism is “a philosophy and social movement which applies evidence and reason to working out the most effective ways to improve the world” (2015, 4-5). Effective altruism’s identity as both a philosophy, or, more precisely, a set of core philosophical commitments, and a social movement places effective altruists in a position in which they must consider the question of which philosophical commitments are essential, such that one must embrace them in order to count as an effective altruist, at least in part in the light of the goal of building a robust social movement capable of advancing its aims. On the one hand, the goal of building a social movement provides a strong reason for effective altruists to embrace a set of core commitments that is as ecumenical as possible, so that no one is excluded who might otherwise identify with the movement and contribute to advancing its central aims. On the other hand, there are a number of risks involved in adopting a broadly ecumenical approach to social movement building. Such an approach can, for example, limit a movement’s ability to challenge conventional wisdom by publicly defending controversial claims, including claims the acceptance of which may have helped to bring together and motivate many early members of the movement. More generally, it can undermine the distinctiveness of a movement, making it less clear how it should be distinguished from other sufficiently similar movements, and why it might represent an important and unique way of thinking about and attempting to address the social problems on which it is focused. A broadly ecumenical approach can also make certain forms of

cooperation among members of a movement more difficult to sustain than they would otherwise be, since disagreements about the core values that the movement should embody and aim to promote will run deeper than they otherwise would. In addition, attempts to be ecumenical in defining a social movement's core commitments can make it more likely that the movement will appear to be, or, in some cases, will in fact be, committed to inconsistent or otherwise implausible combinations of views.

These are not the only risks involved in the adoption of a broadly ecumenical approach to social movement building, but they are risks that, it seems to me, are particularly important for effective altruists to consider in thinking about how to characterize the core commitments of their movement. At the same time, to the extent that effective altruists can be ecumenical in articulating their movement's core commitments without taking on (too many and too much of) these risks, there is, I think, clearly good reason for them to do so. In this paper, my aim is to develop a view about how effective altruists should characterize the movement's core commitments, in light of the challenges presented by the goal of social movement building. In pursuing this aim, I assume that the central moral goals of the movement are worth pursuing, and that the movement has the potential to make genuinely valuable contributions to the world.¹ My discussion will highlight some important limits to effective altruism's ability to be ecumenical deriving from the need to maintain a distinctive identity that is tied to central moral commitments that have, from the outset, animated the movement, as well as the need to avoid commitment to inconsistent or otherwise implausible claims or combinations of claims.² Despite these limits, however, I will argue that there are significant ways in which effective altruism's core commitments both can and should be understood in ways that are substantially more ecumenical than some critical characterizations of the movement suggest they can be. I will

suggest that effective altruism can be understood in terms of a set of core philosophical commitments that leaves room within the movement for a range of views on matters that are, at least from the perspective of the core commitments, subject to reasonable disagreement.³ In addition, I will note a structural similarity between my account of effective altruism's core commitments and John Rawls's idea of an "overlapping consensus" on core principles of justice in a liberal democratic society (1993, lect. IV), but also highlight a key substantive difference that, in my view, contributes significantly to the appeal of my account.

Previous discussions of effective altruism's core commitments, and of the nature and extent of the movement's ability to be ecumenical, have either largely limited themselves to emphasizing its fairly extensive capacity to be ecumenical in response to a range of critiques (e.g. MacAskill 2019; Pummer and MacAskill forthcoming), or focused on particular ways in which the authors believe the movement is not, or cannot be, particularly ecumenical (e.g. Gabriel 2017; Berg 2018). In contrast, I aim to develop an account of the movement's core philosophical commitments that both reflects the fundamental moral concerns that tend to motivate its proponents, and takes seriously both the advantages and the risks of ecumenicism for social movements. This approach to determining how we should understand the movement's core commitments, in turn, can provide the basis for further reflection on how members should understand the intellectual and practical structure that they have reason to promote internally, in particular with regard to the nature and substance of engagement on moral questions. These are extremely important issues for any social movement, but have thus far received little attention in academic work on effective altruism. While I can only offer some relatively brief initial suggestions in this paper, my hope is that my discussion will encourage further thought about the role that moral engagement on relevant matters of disagreement should play within social

movements (including, but not limited to, effective altruism), and what such engagement should look like.

I will proceed as follows in the remainder of the paper. First, in section 2 I will highlight what seem to me to be the central moral concerns that animate the effective altruism movement, and offer an initial characterization of its core philosophical commitments on the basis of these central moral concerns. In section 3, I will note a number of areas of moral disagreement, including several that have motivated critical responses to effective altruism, and consider the extent to which, and the ways in which, effective altruists can and cannot be ecumenical with respect to these issues. I will conclude, in section 4, by describing the way in which, on my view, the core commitments of effective altruism can be thought of as representing a consensus around which the movement should be structured, and note an important difference between the way in which I think members should approach engagement with those within the movement with whom they disagree and the way in which Rawls imagines those who hold incompatible “comprehensive doctrines” (1993, lect. IV) engaging in social cooperation with each other in a just liberal democratic society.

2. Effective Altruism’s Central Moral Concerns

Effective altruists tend to do, and to encourage others to do, things such as:

- Donating money to charitable organizations recommended by effective altruist charity evaluators such as GiveWell and Animal Charity Evaluators.
- Living a modest lifestyle so as to limit one’s negative impact on the world and increase the amount of one’s resources directed to improving the world (Singer 2015, chap. 3).

- Choosing a career on the basis of the capacity of the work itself and/or the earnings from the work to contribute to improving the world in morally important respects (Singer 2015, chap. 4-5; MacAskill 2014; MacAskill 2015, chap. 9).
- Contributing to efforts to determine what is achieved by different approaches to improving the world, so as to increase our ability to make informed choices about where to direct our time and money.
- Making decisions about where to direct the resources and time that one allocates to trying to improve the world on the basis of the best available evidence about what is achieved by different efforts.

Effective altruists are, on the other hand, critical of behavior such as:

- Spending significant amounts of money on unnecessary luxury goods and services for oneself or one's loved ones and friends.
- Donating money to, for example, wealthy universities (Singer 2015, 10-11), one's local opera house, or charities that are, on any plausible view, clearly much less cost effective than others with respect to improving the world.
- Donating money to causes that one happens to care particularly about or feel a special connection to, rather than to others where the donations would, on any plausible view, do significantly more to improve the world (MacAskill 2015, 41-42).
- Donating money to charitable organizations without looking into the available evidence about which organizations improve the world most cost effectively.

The organizations that effective altruist charity evaluator GiveWell recommends focus on providing anti-malarial bednets or preventive drugs, supporting deworming treatments, and providing direct cash transfers to very poor people.⁴ These organizations aim to improve people's lives by helping them avoid contracting malaria, treating intestinal worms that reduce the quality of life of those infected in various ways, and providing cash that beneficiaries can use to improve their lives in ways that they themselves choose. All of these efforts are, it seems to me, best thought of as worth supporting because of their potential to provide for morally important needs of beneficiaries, improve their welfare or quality of life, or, more broadly, to advance their morally important interests. Furthermore, these are the kinds of reasons that effective altruists will, at least most often, offer for supporting the recommended organizations. Effective altruists believe, then, that the needs and interests of individuals provide us with strong reasons to act so as to satisfy those needs and advance morally important interests. This commitment to what I will call *Strong Interest-Based Reasons* for actions such as charitable giving is, in my view, a core commitment of effective altruism.

The organizations recommended by GiveWell focus primarily on improving the lives of people in some of the poorest parts of the world. This focus reflects both the fact that the needs that can be met in these places through the contributions that effective altruists encourage are, on any plausible view, of great moral importance, and the fact that these needs can typically be met more cheaply than any similarly important needs in wealthier parts of the world. Effective altruism has, then, adopted a global orientation, involving a commitment to consider the needs and interests of everyone as equally morally important, regardless of where they happen to live. This *Cosmopolitan Impartiality*, as I will call it, is, in my view, also a core commitment of effective altruism. Those who accept Cosmopolitan Impartiality believe that there are strong

reasons to direct one's charitable resources wherever they will do the most good, impartially understood. Those who reject Cosmopolitan Impartiality, and believe, for example, that they have at least as much reason to direct all of their charitable donations to organizations that benefit their fellow citizens, or others with whom they share particular ties, as they have to direct those donations to organizations that provide aid to foreigners and, impartially speaking, clearly do much more good, are not effective altruists.

The effective altruist charity evaluators determine which charitable organizations to recommend in large part by examining the available evidence regarding what is achieved by different organizations, and how much money is spent in order to achieve it. There are many difficult questions that can be asked about how to interpret the evidence that is available, which metrics ought to be employed in determining which organizations achieve morally important results most cost-effectively, and how to think about the reasons to direct resources to efforts for which evidence of cost-effectiveness is not, and in at least some cases could not, be available (for example, efforts to bring about significant changes in the way that our political or economic institutions are structured). Nevertheless, effective altruists are committed to the view that we should gather as much evidence as we (cost-effectively) can regarding what can be achieved through various efforts to improve the world, and, to the extent that we can, decide what to do in our own efforts to improve the world on the basis of that evidence. This commitment to what I will call *Evidence-Based Decision Making* regarding our efforts to improve the world is, in my view, also a core commitment of effective altruism.

My discussion thus far suggests that effective altruism can, at least initially, be plausibly characterized by core commitments to:

1. Strong Interest-Based Reasons
2. Cosmopolitan Impartiality
3. Evidence-Based Decision Making

These core commitments clearly rule out counting many people as effective altruists. In particular, there are many people who reject Cosmopolitan Impartiality, and, I suspect, a fair number who reject Evidence-Based Decision Making.⁵ Nevertheless, these core commitments at least appear to be consistent with a wide range of more specific moral commitments that are the subject of disagreement both among philosophers and more generally. In the next section I will consider to what extent this initial appearance can be vindicated by considering whether effective altruists face pressure from their core commitments, or for other reasons, to adopt particular positions on additional matters of philosophical controversy, including several that have been the focus of some critical discussions of the movement.

3. How Ecumenical Can Effective Altruism Be?

It is sometimes suggested that effective altruism cannot be ecumenical to any significant extent because it is, in effect, just utilitarianism under another name. It is, on the one hand, easy to see why some may have that impression. The movement can, to a significant extent, be traced to Peter Singer's famous paper "Famine, Affluence, and Morality" (1972). Singer is also a prominent member of the movement, and has long been a committed utilitarian. Two of the movement's most prominent younger philosophers, Will MacAskill and Toby Ord, are also known to have at least broadly consequentialist sympathies in ethical theory. In addition, a recent survey of self-identified effective altruists found that 56% identified as utilitarians, with another

13% identifying as non-utilitarian consequentialists.⁶ Only 8% of respondents identified as deontologists or virtue ethicists, with virtue ethicists outnumbering deontologists.⁷ Although this survey is now several years old, and the movement has grown and evolved in the intervening period, it is at least plausible that the movement is, even now, populated largely by utilitarians and those with more broadly consequentialist sympathies in ethical theory.⁸

Nevertheless, it seems to me clearly a mistake to view effective altruism as simply utilitarianism by another name. After all, even though the numbers are relatively small, there are self-identified deontologists and virtue ethicists who also identify as effective altruists, and there is nothing in either the core commitments that I have identified, or in the recommendations and criticisms that I have identified as typically made by effective altruists, that seems clearly inconsistent with all non-consequentialist ethical theories. In addition, although Singer acknowledges that most effective altruists accept the utilitarian view that values such as freedom and equality are only instrumentally good, insofar as they “are essential for the building of communities in which people can live better lives,” he goes on to say that “[n]o doubt some effective altruists hold that these values are also good for their own sake” (2015, 8-9). It is also useful to note that Singer’s argument in “Famine, Affluence, and Morality” does not depend in any way on consequentialist commitments in ethical theory (though it is, of course, consistent with such commitments). To the extent that the argument in that paper (as opposed to Singer’s own broader philosophical commitments) represents a common and important reason that people identify as effective altruists, it is misleading to think that the movement is defined by explicitly utilitarian commitments.

If effective altruism cannot be identified with utilitarianism, or even with consequentialism more broadly, then effective altruists should be able to pursue their aim of

building a robust social movement by appealing to individuals with a broader range of theoretical commitments in ethics, as well as those with no particular commitments at the level of moral theory. In the remainder of this section I will highlight a number of areas of widespread moral disagreement, and consider which theoretical questions effective altruists can, at least qua effective altruists, consider to be open, and on which they must, or are at least under significant theoretical pressure to, adopt particular contested positions.

3.1. Rights and Deontological Constraints

Consequentialist moral theories have well-known difficulties fully accommodating widespread intuitions about rights and deontological constraints. For example, it is difficult to see how a consequentialist can deny that an agent ought to push one person off a bridge if doing so would prevent a runaway trolley from killing five other people,⁹ despite the fact that most people have the intuition that doing this would be wrong.¹⁰ If this intuition is correct, then it seems likely that the best explanation for its correctness is that pushing the person off the bridge would violate a right that takes priority over bringing about the best outcome, or violates a deontological constraint that cannot be violated even in (most) cases in which violating it would bring about the best outcome. Is holding that there are rights and/or deontological constraints incompatible with what I have suggested are the core commitments of effective altruism, or otherwise suspect from the perspective of effective altruism?

It does not appear that there is any inconsistency here. It seems entirely possible to believe, for example, that when it comes to deciding where to direct one's charitable resources, one should take everyone's interests to matter equally, regardless of whether they are, for example, one's fellow citizens or not, and to also believe that one ought not, for example, steal

resources from one's wealthy neighbor in order to have more to donate to much needier people elsewhere.¹¹ The same seems true with respect to the effective altruist commitments to Strong Interest-Based Reasons and Evidence-Based Decision Making. In addition, nothing in the recommendations of effective altruist charity evaluators, or in the kinds of criticisms that effective altruists often make of common spending and donation practices, seems inconsistent with accepting rights-based or deontological constraints on promoting the good.¹²

On reflection, the compatibility of effective altruism and a commitment to rights and/or deontological constraints should not be especially surprising. The focus of the effective altruism movement is on beneficence, and the moral reasons that support it. In other words, effective altruism consists in a set of core commitments regarding the kinds of reasons that we have to act so as to benefit others, and, as I will suggest later in the paper, the obligation-generating force of those reasons. While it seems to me to be the case that consequentialists tend, on average, to accept more stringent accounts of our reasons of beneficence and, relatedly, of our obligations of beneficence than those who reject consequentialism, there is no inconsistency involved in accepting that there are rights and/or deontological constraints that individuals must respect, and also accepting that we have very strong moral reasons to aid others when we can do so without violating those rights or constraints. Effective altruists, then, need not adopt any particular contested position regarding rights and/or deontological restrictions. On this issue, effective altruism can be quite ecumenical.

3.2. Justice and Distributive Principles

Utilitarians believe that we have most moral reason to direct our charitable resources in whatever way will maximize aggregate utility. On this view, we will sometimes have reason to

direct an entire pool of resources to one needy group, and none to another, despite the fact that the groups are equally badly off (Gabriel 2017, 459-60). This will be true if, for example, each dollar would produce a slightly greater welfare benefit if directed to the first group than it would if directed to the second. Utilitarianism also implies that we will sometimes have reason to direct our charitable resources to better off people rather than to worse off people (Gabriel 2017, 460-61). This will be true if the better off people would gain more welfare from each dollar directed toward helping them than the worse off people would. Many people, however, have intuitions that suggest that maximizing aggregate utility is not the only thing that matters, morally speaking, in cases of these kinds. They think, for example, that justice can require that resources, including aid resources, be distributed in a broadly egalitarian manner, or that there are reasons to prioritize the needs and interests of those who are worse off than others.¹³ Are effective altruists committed to rejecting these views?

Iason Gabriel claims that there are both “thick and thin versions of effective altruism” (2017, 458). The thick version, on his view, involves a commitment to utilitarianism, while the thin version merely involves a commitment to “using a substantial amount of o[ne’s] spare resources to make the world a better place,” and is “compatible with a wide range of moral theories” (Gabriel 2017, 458). If, however, those who accept what Gabriel calls the thin version, but reject the thick version, are nonetheless effective altruists, then it is misleading to describe the thick version as a version of effective altruism. Instead, effective altruism just is the thin version, and it so happens that many effective altruists (that is, many who accept the claims that constitute the thin version) also believe that utilitarianism is the correct moral theory.

If it is understood in terms of the claims that constitute Gabriel’s thin version, or in terms of the core commitments to Strong Interest-Based Reasons, Cosmopolitan Impartiality, and

Evidence-Based Decision Making that I have described, then it seems clear that there is nothing about effective altruism that is inconsistent with broadly egalitarian distributive principles, or with the view that there are reasons to prioritize the needs and interests of the worst off even when doing so will not maximize aggregate welfare. Strong Interest-Based Reasons and Cosmopolitan Impartiality do together imply that we have strong reasons to advance people's morally important interests as much as possible, taking everyone's interests to matter equally. But this does not rule out thinking, for example, that an equal welfare gain has greater moral importance when it benefits a worse off person than it would have if it benefitted a better off person.¹⁴ From the perspective of the core commitments that I have identified, then, effective altruists can disagree about the distributive principles that ought to guide us in our efforts to improve the world. On this issue too, effective altruism can be ecumenical.

3.3. Political Action and Institutional Change

Perhaps the most common criticism of effective altruism is that it is insufficiently focused on political action and the importance of achieving large-scale global institutional change.¹⁵ The central concern seems to be that the movement's commitment to recommending that individuals direct their time and resources toward whichever efforts appear, given the available evidence, likely to do the most good, will make whatever commitments its members have to working for institutional change objectionably contingent, since it will always be an open question whether political action appears, for any particular individual deciding what to do with her time and resources, to offer the best prospects for improving the world. Proponents of this criticism believe that our commitment to directing our time and resources toward efforts to bring about institutional change should not be contingent in this way.

One interesting thing to note about this criticism of effective altruism is that, unlike those that suggest that it cannot accommodate widely held non-utilitarian views about rights, deontological constraints, or distributive principles, it claims that effective altruism is objectionable because it is, in at least one way, *too ecumenical*. Those who advance this criticism claim that effective altruism is objectionable because it lacks a core commitment to the pursuit of certain kinds of global institutional changes, typically those advocated by people who identify with the political left. On their view, any social movement aimed at addressing issues such as global poverty ought to be explicitly structured around working for particular kinds of global institutional changes.

The effective altruist commitments to Strong Interest-Based Reasons and Evidence-Based Decision Making seem to clearly rule out the kind of core commitment to structuring the movement around efforts to bring about the particular kinds of global institutional reforms advocated by the critics. Even if there were good reason to believe that the kinds of institutional reforms that they advocate would, if achieved, greatly benefit the global poor and thereby improve the world,¹⁶ if the prospects of achieving them seem low enough, then effective altruists will not advocate devoting time and resources to efforts to bring them about.

One reason that the prospects of achieving certain global institutional reforms might be low is that there is a great deal of disagreement about, for example, the appropriate role of governments in the economy. Some people, for example at least some libertarians, might accept that there are strong reasons for individuals to direct resources to, for example, effective efforts to improve the lives of the global poor, while also believing that there are deontological constraints that make it impermissible for governments to advance similar efforts via coercive policy.¹⁷ All else equal, it is a good thing from the perspective of effective altruist social

movement building efforts if such people can be included in the movement, just as it is a good thing that those who accept that there are deontological constraints that apply to individuals can be included.

It is important to note that the fact that effective altruism can, on my view, be at least fairly ecumenical with respect to views about the appropriate role of government does not prevent the movement from including sub-groups that work toward particular policy goals. One feature of any ecumenical social movement will be internal disagreement about issues that are not settled by the movement's core commitments. These disagreements should not, in a well-functioning movement, significantly hinder cooperation around the issues that are settled by the core commitments.¹⁸ In addition, there are certain important policy goals that are likely to be endorsed by virtually anyone who would count as an effective altruist, even on a maximally ecumenical view. Two possible examples are progressive criminal justice reform in the United States¹⁹ and ending agricultural subsidies in wealthy developed countries.

With respect to political commitments, then, effective altruism is, as some critics have suggested, at least fairly ecumenical. It seems to me that, although there are some potential costs to being ecumenical in this way, on the whole it is a good thing for the movement.

3.4. Obligations and Demandingness

Some critics of effective altruism claim that it is objectionable because it involves a commitment to an overly demanding account of our obligations of beneficence.²⁰ In response to this concern, some have argued that effective altruism need not be understood as involving any commitment to obligations of beneficence at all (Bowen 2016; Pummer 2016).²¹ On this view, what is essential to effective altruism is a commitment to the claim that those who do devote

resources to efforts to improve the world have reason, or a *conditional* obligation, to direct those resources in ways that will do the most good. Accepting this claim is consistent, it is suggested, with accepting that there is nothing wrong with devoting no resources whatsoever to efforts to improve the world. And accepting both of these claims is, on this view, consistent with the core commitments of effective altruism.

Dillon Bowen (2016) suggests that effective altruism should be understood as a corrective for the cognitive biases that lead many people who give in order to try to improve the world to give ineffectively, that is, to give to efforts that do less good than other efforts to which they could have given instead. He claims that at least most people would, on reflection, prefer to give to more effective rather than less effective efforts, so that the cognitive biases to which we are subject tend to lead us to act contrary to our own preferences (Bowen 2016, 81, 88-89). And this inconsistency between how many people in fact donate and the preferences that those same people would have, on reflection, about where to donate is, on his view, the problem that effective altruism should be understood as seeking to remedy: “[i]neffective altruism...is problematic in the sense that it violates our preferences. When people behave as...ineffective altruists, I do not necessarily think they are violating a moral duty so much as behaving...in a way they themselves would probably disapprove of in light of their own values” (Bowen 2016, 88).

Theron Pummer (2016) defends an alternative view, according to which it can be morally wrong to devote resources to a less effective effort to improve the world, rather than to a more effective effort that one could have supported instead, even if it would not be wrong to refrain from giving anything at all. In other words, he defends the claim that we can have an obligation to give to the most effective effort that is conditional on the fact that we are going to give to

some effort or another. This condition obligation is captured, Pummer claims, by the following principle:

Avoid Gratuitous Worseness (weak): It is wrong to perform an act that is *much worse* than another, if it is *no* costlier to you to perform the better act, and if all other things are equal (Pummer 2016, 84, emphasis in original).²²

Pummer acknowledges that if we accept that there is no unconditional obligation to give to efforts to improve the world, then the same considerations that ground the permission not to give will also, in at least many cases, ground permissions to give to less effective efforts rather than more effective ones (2016, 82, 92). For example, if I am permitted to refrain from giving in order to spend all of my disposable income on a personal project that I care about, such as improving my golf game, then I will also be permitted to give (some portion of) my disposable income to a charity that does work that I happen to particularly care about, rather than giving it to another one that does more good, impartially speaking. This is because we cannot plausibly count making sacrifices with respect to the pursuit of my project of improving my golf game as a cost to me, but deny that giving up donating to a charity whose work I happen to care a great deal about is also a cost. Pummer's principle, then, counts as wrong only those donations to less effective rather than more effective charities that are made by people who have no particular attachment to, or caring attitude toward, the work done by the less effective charities.

Pummer's principle would allow effective altruists to criticize a certain kind of charitable giving that is at least not uncommon. Many people give to organizations whose work they do not care about any more than they care about the work of other, perhaps more effective

organizations. They might, for example, give as a result of seeing a particularly well-made ad on television, or because a friend happened to mention giving to the organization. Insofar as these donors could just as easily have given instead to a different, more effective organization, Pummer's principle implies that they acted wrongly. Importantly, however, donations of this kind seem to be the *only* ones that the principle will rule out. And this should seem clearly inadequate to effective altruists, since the movement is clearly committed to advocating for much broader changes to individuals' charitable giving practices.²³

Both Bowen and Pummer's accounts of how effective altruism can avoid commitment to (unconditional) obligations of beneficence strike me as quite problematic, both from the perspective of what seem to be among the movement's core philosophical commitments, and from the perspective of social movement building. Notice, for example, that both accounts deny that there is a basis within effective altruism's core commitments for making many of the kinds of apparently moral criticisms that effective altruists in fact tend to make. If, for example, the only thing that can be problematic about donations to less effective efforts to improve the world, from an effective altruist perspective, is that those making them would actually prefer to donate in more effective ways, then there are no effective altruist-based grounds for criticizing those who genuinely prefer to benefit wealthy universities or local opera houses rather than the global poor. And if, from the perspective of effective altruism's core commitments, there are no unconditional obligations to contribute to efforts to improve the world, then there are no effective altruist-based grounds for criticizing those among the well off who spend all of their disposable income on luxury goods and services for themselves. Since effective altruists do tend to criticize those who appear to have a genuine preference for supporting causes that are, from an impartial perspective, much less important than others that they could support, as well as those

who give nothing at all to efforts to improve the world, there is at least some reason to be concerned about accounts of the movement's core commitments that provide no grounds for such criticisms.

More importantly, however, both Bowen and Pummer's accounts are inconsistent with accepting Cosmopolitan Impartiality as a core moral commitment of effective altruism. This is because they provide no grounds for criticizing those who reject that commitment entirely, either by always prioritizing their own interests, or because they happen to genuinely care more about, for example, benefitting fellow citizens than benefitting the global poor, so that all of their donations go toward local efforts rather than much more effective global ones. Cosmopolitan Impartiality, however, seems clearly to lie at the core of effective altruism. Effective altruists believe that there is at least some sense in which, and some degree to which, we all ought to be committed to the equal moral importance of everyone's needs and interests. And this commitment will entail that (at least) the well off have at least some unconditional obligations to contribute to efforts to improve the world.

Accepting that effective altruists must endorse unconditional obligations to contribute to efforts to improve the world does not, however, commit us to any particular view about how demanding those obligations are. On the issue of demandingness, effective altruism can, it seems to me, be at least fairly ecumenical. Strong Interest-Based Reasons and Cosmopolitan Impartiality together require effective altruists to accept that a commitment to the equal moral importance of everyone's needs and interests must play a non-trivial role in guiding our behavior, and in particular our contributions to efforts to improve the world. But they do not commit us to any particular view about how this commitment can or cannot be balanced against other commitments that individuals might also permissibly have, such as commitments to

personal projects and special relationships. There is, then, room within the movement for disagreement about the demandingness of our obligations of beneficence.

Despite the fact that there is room for a certain amount of disagreement about how demanding our unconditional obligations of beneficence are, there are limits to how ecumenical effective altruism can be on this issue. It is difficult to be very precise about what these limits are. Intuitively, however, a multi-millionaire who gives one tenth of one percent of her income away, but is careful to ensure that it is directed to very effective organizations that improve the lives of badly off people, cannot plausibly claim to be an effective altruist. Rather, it would seem that effective altruism at least requires well off people to give up enough in the way of resources that they are making some non-negligible sacrifice, relative to an alternative in which they do not give at all.

Indeed, even this standard seems, on reflection, too weak. This is because a commitment to effective altruism requires, as I have argued, a commitment to Cosmopolitan Impartiality. And although maintaining a commitment to Cosmopolitan Impartiality does not, it seems to me, require that well off agents always act in the ways that are required by the most demanding moral theories, it does require that they are willing to make some significant sacrifices in order to benefit those who are much worse off. Though I cannot offer a complete defense here, in my view the best way to conceive of the floor for what effective altruism requires here is that individuals must integrate a serious commitment to the equal moral importance of everyone's needs and interests into the overall set of commitments that informs their behavior (Berkey 2016). On this view, determining how much each individual is required to donate to effective organizations requires determining how much she would have to give, taking into account the resources available to her and other relevant features of her life, in order to be able to plausibly

claim to have integrated a serious commitment to the equal moral importance of everyone's needs and interests into her overall set of commitments. This will, I think, require well off people to give much more than a trivial amount. It is, however, at least debatable just how extensive the sacrifices required by this standard are.²⁴

Some might think that even if I am correct that effective altruism is, philosophically speaking, best understood as involving a commitment to unconditional obligations of beneficence, from the perspective of social movement building it is problematic to insist on a commitment to such obligations.²⁵ We might worry, for example, that recruiting people to the movement will be made more difficult because people do not like having it suggested to them that they have moral obligations that they are currently failing to meet.

While it may be true that some who might otherwise be sympathetic to the movement would be deterred by its commitment to unconditional obligations of beneficence, it seems to me that the costs of rejecting that commitment would be much higher. First, the movement would lose the ability to coherently make many of the criticisms of common spending and donating habits that have helped it gain prominence as a (collective) critic of common sense morality. Without the ability to claim that everyone who is not already an effective altruist has moral reasons, and indeed obligations, to change their spending and donation habits, the movement would lack an important means of attempting to recruit those who are not already involved, namely moral persuasion. And in the absence of a commitment to engaging in moral persuasion efforts, it would, I think, be unclear why effective altruism should be thought of as a *social movement* at all, rather than as merely a group of people, like many others, committed to doing good things that they just happen to care about more than other good things that they could equally well be doing instead. It is acceptance of the moral claim that *everyone* ought to be

contributing to effective altruist efforts, or at least ought to endorse such efforts and the core commitments that support them, that, in my view, makes sense of the concern to build a social movement.

There is, it seems to me, something puzzling in the thought that a person might be committed to a social *movement* (as opposed to, say, a social club or a particular charitable cause) while believing that endorsement of and commitment to acting on the values that animate the movement is, morally speaking, optional.²⁶ It is difficult to imagine that this was the attitude of anyone committed to, for example, the civil rights movement, or the women's liberation movement.²⁷ Indeed, I suspect that at least many who contributed to these movements were motivated to do so as a result of becoming convinced that supporting them was morally obligatory.²⁸ Far from being off-putting, for at least some people convincing arguments for the claim that there are moral obligations to contribute to the goals of a social movement can be extremely motivating, in ways that other kinds of recruitment efforts are unlikely to be. I do not think, then, that the goal of building a social movement provides any reason to reject unconditional obligations of beneficence as a core commitment of effective altruism.

3.5. Permissible Partiality and Special Obligations

Consequentialist moral theories are often criticized as incompatible with powerful intuitions about permissible partiality and special obligations to particular people, such as family members and friends. Are the core commitments of effective altruism, including now the commitment to unconditional obligations of beneficence, incompatible with those intuitions?

While the commitment to unconditional obligations of beneficence does imply that the well off cannot permissibly devote *all* of their resources to loved ones and friends, it does not

necessarily rule out the permissibility of devoting a substantial portion to them. It also does not rule out special obligations to devote some portion of the resources that one is not obligated to dedicate to satisfying obligations of beneficence to particular people, such as one's children. The core commitments of effective altruism entail that there is some portion of a well off person's resources that she is obligated to devote to the most effective efforts to improve the world to which she can contribute. For at least most well off people, I have argued that what is required will be much more than trivial, in virtue of the requirement to integrate a serious commitment to the equal importance of everyone's needs and interests into the overall set of commitments that informs their behavior. The core commitments, however, do not by themselves settle the question of *exactly how much* of a well off person's resources she is obligated to devote to such efforts. It is at least open to effective altruists to argue that well off people are not obligated to devote an *extremely large* portion of their resources to efforts to improve the world. The movement can, therefore, be at least fairly ecumenical with regard to views about permissible partiality and special obligations.

3.6. *Non-Human Animal Interests*

Many effective altruists believe that the most effective way that well off people can improve the world is to contribute resources to efforts to reduce the suffering of non-human animals, and in particular farmed animals.²⁹ At least most who hold this view believe that the interests of non-human animals are no less morally important than the equivalent interests of humans.³⁰ To what extent must effective altruists embrace this view, and the implications for how we ought to direct our efforts to improve the world that many of its proponents take it to have?

Because there is now a great deal of evidence that many non-human animals possess at least fairly sophisticated cognitive capacities (Varner 2012, chap. 5-8), in addition to the basic capacities to experience pleasure and pain, the commitments to Strong Interest-Based Reasons and Evidence-Based Decision-Making together put substantial pressure on effective altruists to accept not only that non-human animals have the kind of moral status that makes it the case that we are obligated to take their interests into account in determining what we ought to do, but also that many of the ways that certain animals, and in particular farmed animals, are treated are seriously morally problematic. It is difficult to see how one could accept the core commitments of effective altruism that I have described and nonetheless hold that there is nothing especially problematic about, for example, the ways that animals are typically treated in factory farms. Factory farmed animals seem clearly to share enough of the interests that effective altruists take to ground our reasons to prevent human suffering to ensure that effective altruists are committed to accepting that there are strong reasons to be concerned about the ways that our practices affect those animals.

Despite this, there is, it seems to me, at least some room for reasonable disagreement among effective altruists about, for example, whether the equivalent interests of humans and non-human animals must count as equally morally important, and, relatedly, about whether working to reduce the suffering of farmed animals is the most effective way that at least many of us can improve the world. There is, it seems to me, not necessarily anything inconsistent with the core commitments of effective altruism in the view that because non-human animals lack at least some of the most sophisticated cognitive capacities that typical humans possess, they have somewhat lesser moral status than typical humans (though they still, of course, have significant moral status). Those who accept this view might also accept that, for example, an equivalent

amount of suffering matters somewhat less when it is endured by a being with lesser moral status than it would if it were endured by a being with greater moral status.

Although there is a range of views that might be adopted by those who believe that some beings have lesser moral status than others about *how much less* the suffering of beings with lesser moral status matters in comparison with the suffering of beings with greater moral status, it seems to me that the core commitments of effective altruism rule out thinking that this difference can be more than modest. The overlap in the central interests that humans generally share with at least many non-human animals requires that Strong Interest-Based reasons be understood in a way that implies that the interests of such animals provide reasons that are least similar in strength to those provided by equivalent human interests. In addition, Cosmopolitan Impartiality rules out any additional discounting of the weight of these reasons for particular agents, since this would be objectionable for the same kind of reason that discounting the weight of reasons provided by the interests of, for example, distant humans is. If we are not permitted to count the interests of humans for less in virtue of the fact that they are distant strangers rather than, for example, co-citizens, then we are not permitted to count the interests of non-human animals for less simply in virtue of the fact that they are members of different species than us.³¹

Because of this, there is only a limited range of views about the relative moral importance of, on the one hand, reducing the suffering of farmed animals, and, on the other, potentially competing priorities such as improving the lives of the global poor, that might be held by effective altruists. Because of the empirical difficulties involved in comparing the significance of the interests that are at stake, however, there will be much more reasonable disagreement about which efforts effective altruists should prioritize in the actual world, even within the more limited theoretical disagreement that is consistent with the core commitments.

Effective altruism can, then, be fairly ecumenical (though, I have suggested, far from maximally so) about the relative moral importance of non-human animal interests and the equivalent interests of humans, as well as about whether the cause of reducing non-human animal suffering should, given the current state of the world, be prioritized over other morally important causes when it comes to resource allocation decisions.³²

3.7. Future Generations and Existential Risk

One of the most controversial positions that some effective altruists have embraced is that the most effective way that we can use current resources to improve the world is to invest in mitigating “existential risks,” that is, risks, including very low probability risks, that threaten to either cause humanity to go extinct or to “permanently and drastically curtail its potential” (Bostrom 2002; see also Bostrom 2013). Examples of existential risks include the possibility that a large asteroid will collide with the Earth, the possibility that humanity will destroy itself via a nuclear holocaust, and, perhaps most controversially, the possibility that a poorly programmed extremely powerful artificial intelligence will annihilate humanity (Bostrom 2002).

Those who hold that there are strong reasons to prioritize mitigating existential risks believe that a massive amount of morally relevant value would be lost if an extinction-causing event were to occur, and that this fact is what explains why it is more important to mitigate such risks than, for example, to do more to alleviate global poverty in the near term.³³ The morally relevant value at stake, according to proponents of mitigating existential risks, consists in all of the value that would be contained in the lives worth living that would be prevented from coming into existence by an extinction-causing event. Nick Bostrom argues that even given an especially conservative estimate of the number of life years that could be lived in the future if we avoid

extinction-causing events, the value of reducing existential risk by one millionth of one percent should be considered as valuable as saving a million lives (Bostrom 2013, 18-19).

Those who believe that there are strong moral reasons to mitigate existential risks, and that resources would do more good if directed toward that aim, rather than toward other morally significant aims such as alleviating global poverty or reducing the suffering of factory farmed animals, hold this view because they accept that a world in which more people come into existence and live good lives is, all else equal, better than a world in which fewer people do so. In addition, they take the fact that a world in which more people come to exist and live good lives is better to provide us with moral reasons to do what we can to ensure that more people rather than fewer come to exist and live good lives (Bostrom 2013, 18-19; Parfit 1984, 453-54).

While the view that we have strong reasons to do what we can to ensure that more people rather than fewer come to exist and live good lives is plausible, it is also quite controversial. Many procreative and population ethicists, for example, believe that there are no moral reasons to bring into existence people who would have good lives, since failure to bring such people into existence cannot wrong them (after all, they don't exist).³⁴ If this view is correct, then the value that would be contained in the lives of future people if we avoid an extinction-causing event would not provide us with strong reasons to invest in mitigating existential risks.

There is, it seems to me, nothing in the core commitments of effective altruism that requires effective altruists to adopt any particular position on the highly controversial and difficult issues in procreative and population ethics that are relevant to the debate about the moral value of mitigating existential risk. It is consistent with those core commitments to think either that mitigating existential risk is the most important thing to which we can be devoting

resources, or that there is no reason to devote resources to these efforts rather than to efforts to, for example, alleviate global poverty.

Both Strong Interest-Based Reasons and Cosmopolitan Impartiality do, however, seem to me to generate significant pressure on effective altruists to consider the interests of future people who will exist regardless of what we do to count for no less than the interests of currently existing people. This requirement goes well beyond the requirement to ensure that such people are protected from serious threats. The latter requirement is relatively uncontroversial, even among those who do not consider themselves effective altruists. The requirement to count the interests of future people who will exist regardless of what we do for no less than the interests of present people, however, implies that fairly radical changes in the ways that most people tend to approach conflicts between the interests of existing people and the interests of those who do not yet, but will exist, must be made.

There are, then, some significant limits to how ecumenical effective altruism can be with regard to controversial questions about the nature and grounds of our obligations to future generations. Importantly, however, with regard to the particular issues having to do with existential risk, the movement can be quite ecumenical.

4. The Structure of the Movement

My discussion thus far provides sufficient grounds for thinking that effective altruism can be at least fairly ecumenical with regard to a number of issues about which there is substantial moral disagreement. At the same time, it, like any social movement, must have a core set of commitments around which it is structured. I have suggested that, in addition to the three core commitments described in section I (Strong Interest-Based Reasons, Cosmopolitan Impartiality,

and Evidence-Based Decision Making), effective altruism should be understood as having a core commitment to the claim that there are at least some unconditional obligations of beneficence.

These four commitments can be thought of as the philosophical core of effective altruism. Those who accept them should consider themselves effective altruists. And since they are all quite plausible, and accepting them is consistent with a range of views about a number of other disputed moral issues, the prospects for building a social movement around them seem at least potentially promising.

I suggest that acceptance of the four core commitments that I have described be thought of as playing a role within the effective altruism movement that is at least somewhat structurally similar to the role that Rawls takes commitment to core principles of justice to play within a just liberal democratic society.³⁵ For Rawls, disagreement about a wide range of moral questions is inevitable even in a fully just society (1993, xviii-xix). Nevertheless, he claims that the goal of maintaining social cooperation on fair terms can be achieved so long as all members of society are committed to certain core liberal ideas, such as the free and equal status of citizens and the importance of protecting certain basic liberties. Commitment to the “overlapping consensus” represented by these core liberal ideas is, for Rawls, sufficient for being a liberal, in the relevant sense, even if one rejects other views that most liberals endorse.

My suggestion is that the central moral goals of effective altruism can be advanced by a movement structured around the core commitments that I have described, despite ongoing deep disagreements among members about a range of other moral questions. If this is correct, then it gives effective altruists a reason to endorse the broadly ecumenical approach to building their social movement that I have suggested, while at the same time insisting on the core commitments that I have identified.

Let me conclude by noting one important difference between the kind of engagement among members who disagree about broader moral issues that I think the effective altruism movement should seek to foster, and Rawls's view about the kind of engagement that is necessary, as a matter of justice, in a liberal democratic society. For Rawls, citizens of a just liberal democratic society do not have reasons deriving from their common goal of maintaining social cooperation on fair terms to attempt to resolve their broader moral disagreements by engaging in ongoing debate about the issues. That is, achieving the goals that they have, qua citizens of a just liberal democratic society, does not require engagement with the aim of developing better justified views about matters on which there is deep moral disagreement, and so it is not essential that liberal societies encourage such engagement.

The effective altruism movement, on the other hand, has as a central moral goal improving the world in morally important respects. From the perspective of this central moral goal, it is important to develop views that are as well justified as possible about exactly what count as morally important improvements to the world, and which improvements are more important than others. Since these are issues about which effective altruists will reasonably disagree, members will have strong reasons to foster a commitment within the movement to robust engagement with the aim of making progress thinking through the issues.

The recognition that there is reasonable disagreement about a range of moral questions among those who accept the core commitments that I have described should, given the effective altruist goal of improving the world as much as possible, also lead members of the movement to think it particularly important to bring in people who will challenge the most widely held views among existing members that fall outside of the core commitments. Effective altruism, then, is a social movement that should embrace a conception of itself as having a strong interest in

generating a great deal of internal moral debate. This, it seems to me, is a striking and deeply appealing feature of a social movement aimed at improving the world.

I am grateful to audiences at the 2018 Great Lakes Philosophy Conference, the 2018 Braga Meeting on Ethics and Political Philosophy, the 2018 Society for Applied Philosophy Conference, the Rhode Island College Department of Philosophy, and the Lilly Family School of Philanthropy at Indiana University-Purdue University, Indianapolis. Rob Hughes, Julian Jonker, Alida Liberman, Carson Young, and David Zaring provided helpful written comments. I have also benefitted from discussions with Eilidh Beaton, Amy Berg, Matt Braich, Stephanie Collins, Matt Duncan, Sarah Light, Catriona McKinnon, Andras Miklos, Darrel Moellendorf, Amy Sepinwall, Patrick Taylor Smith, Aaron Smuts, and Marty Sulek.

Notes

¹ I believe that these claims are true, but will not defend them in this paper.

² I will assume that there are, at least in most cases, strong reasons for genuinely valuable social movements to avoid publicly embracing moral claims that are inconsistent or, on reflection, rather implausible. It may be that there are unusual circumstances in which publicly embracing such moral claims would, on the whole, help such a movement recruit members and advance its goals. But in most cases opponents of a movement will be able to undermine it by highlighting its commitments to inconsistent or implausible views. In addition, there may be moral reasons for members of a social movement to avoid publicly endorsing inconsistent or implausible views even if doing would not risk undermining the movement's central efforts.

³ My arguments in this paper expand upon my earlier discussions of the core commitments of effective altruism in Berkey (2018, 146-48; 2019a).

⁴ GiveWell's charity recommendations can be found here: <https://www.givewell.org/charities/top-charities>.

⁵ In addition, surely many people simply have not thought very much about these claims, and do not have particular views about them. These people do not fail to count as effective altruists in virtue of rejecting the claims, but merely in virtue of failing to endorse them.

⁶ The survey can be found here: <https://eahub.org/sites/eahub.org/files/SurveyReport2015.pdf>. It is unclear whether the survey itself provided subjects with an account of the distinction between utilitarianism and consequentialism, or whether respondents simply answered the question on the basis of their own understanding of the distinction. If the latter, then I suspect that those who identified as non-utilitarian consequentialists did so primarily because they reject hedonistic and preference-satisfaction accounts of the good.

⁷ Although the numbers are low, it is, I think, at least a bit interesting that virtue ethicists outnumber deontologists, since I suspect that deontologists outnumber virtue ethicists among academic philosophers. This suspicion is supported at least a bit by the responses to the recent PhilPapers survey question on philosophers' commitments in normative ethics. That question yielded 25.9% of respondents (241/931) claiming to accept or lean toward a deontological view, and 18.2% (169/931) claiming to accept or lean toward a virtue ethical view. The preliminary results are available at: <https://philpapers.org/surveys/results.pl>. Thanks to (removed for review) for encouraging me to look at the PhilPapers survey results.

-
- ⁸ It is worth noting, however, that 22% of respondents to the survey stated that they either had no particular theoretical commitments in ethics, or were unfamiliar with the competing ethical theories.
- ⁹ This case originally appears in Thomson (1985, 1409-10).
- ¹⁰ Of course, in some cases there may be reasons to believe that the long-term, indirect consequences of doing this would be worse than the consequences of not doing it (perhaps, for example, doing it would weaken social norms against directly harming others to an extent that would lead to greater overall harm in the long run). In addition, some forms of rule-consequentialism might imply that one ought not push the person off the bridge in a case of this kind. I am grateful to an anonymous reviewer for encouraging me to consider these possibilities.
- ¹¹ For a defense of the view that theft for the purpose of donation is, at least in some circumstances, permissible, see Unger (1996, chap. 3). A different defense of the view that theft from the rich in order to give to the poor is at least sometimes permissible is offered in Von Platz (2016).
- ¹² For a slightly more skeptical discussion, see Gabriel (2017, 461-62).
- ¹³ For relevant discussion, see Skelton (2016, 141).
- ¹⁴ For example, it is at least plausible that a distribution in which A has 20 units of welfare and B has 50 is impartially better than a distribution in which A has 15 units and B has 56. And there is nothing in the content of either Strong Interest-Based Reasons or Cosmopolitan Impartiality that would require that effective altruists reject this claim.
- ¹⁵ For examples of this criticism see Lichtenberg (2015); Srinivasan (2015); Herzog (2016). An early articulation of similar concerns can be found in Gomberg (2002). For discussion, see Gabriel (2017, 468-70); Berkey (2018).
- ¹⁶ I am inclined to believe that this is true of some, though not necessarily all, of the global institutional reforms favored by critics of effective altruism on the political left.
- ¹⁷ There are, in fact, a number of people who identify as both libertarians and effective altruists.
- ¹⁸ In practice, of course, disagreements about matters that lie outside of a movement's core commitments do often hinder cooperation around the core commitments. There are, then, in practice always at least some risks involved in building an ecumenical social movement. These risks, however, must be weighed against the benefits of building an ecumenical movement, which can also be substantial.
- ¹⁹ For discussion of effective altruist efforts in this area, see MacAskill (2015, 185-87).
- ²⁰ This objection is, of course, also commonly made against consequentialist moral theories. I discuss how we should understand the objection, and its force against potential individual and collective obligations, respectively, in Berkey (2016) and (2019b).
- ²¹ Also relevant is William MacAskill's recent discussion of how effective altruism should be defined (MacAskill 2019). Since this paper was drafted before MacAskill's discussion appeared, I do not discuss it in detail here. My view, however, is inconsistent with his account of effective altruism as a conjunction of two projects that does not require commitment to any particular normative claims.
- ²² Later in the paper, Pummer suggests that we might consider accepting a stronger version of the principle that requires performing the better act even at some, perhaps substantial, cost to yourself (2016, 93-94). It seems to me, however, that any argument for a stronger version of the principle will also provide support for the kind of unconditional obligation to give, even at some cost to yourself, that Pummer wishes to avoid committing himself to.
- ²³ For a more detailed discussion and critique of Pummer's view, see Berkey (forthcoming).
- ²⁴ I am grateful to an anonymous reviewer for encouraging me to clarify the points made in this and the previous paragraph.
- ²⁵ One way that this claim might be defended is by appealing to Henry Sidgwick's claim that morality can be "esoteric," in the sense that the principles that experts ought to advocate publicly might differ from those that they ought to accept and employ as a guide in their own decision-making; see Sidgwick (1907, 489-92); for discussion see Skelton (2016, 142-43). In my view, even if some esotericism in morality is acceptable, it is difficult to see how a social movement that takes itself to be structured around core philosophical commitments could function in an effective way, and in a way that we

would have reason to find otherwise acceptable, with an “elite” advocating to other members, and prospective members, supposedly core philosophical claims that they themselves reject. One important reason for this is that, as I will discuss in more detail in section 4, effective altruism, as a social movement, has strong reasons to foster an internal culture that encourages a great deal of open debate about challenging ethical issues. And a movement structured by an elite misleading members about their own fundamental commitments could not effectively develop this kind of culture.

- ²⁶ Whether commitment to a social movement is compatible with believing that endorsement of the movement’s core commitments is optional, morally speaking, depends on what account of the defining features of social movements we should endorse. There is relatively little philosophical literature that addresses this question. My suggestion is that because it seems counterintuitive to think that one can be committed to a social movement while believing that those who reject that movement’s core values are not necessarily making any mistake, any plausible account of what a social movement is will imply that all such movements are organized, in some way or other, around a set of core values that members believe ought to be endorsed by everyone. For discussion, see Kolers, (2016, 586-88).
- ²⁷ Of course, many members of those movements were resisting their own oppression, and so had self-interested in addition to moral reasons for their commitment. This does not, however, in any way call into question my suggestion that they almost certainly believed that everyone had a moral obligation to support the movements’ goals.
- ²⁸ I suspect that this is also true of at least many who have come to identify as effective altruists, since it seems likely that a common reason for becoming a member of the movement is that one read Singer’s “Famine, Affluence, and Morality,” and was convinced that there is at least something importantly right about the argument.
- ²⁹ See, for example, the answers to frequently asked questions provided by Animal Charity Evaluators, available here: <https://animalcharityevaluators.org/about/background/faq/>.
- ³⁰ The seminal defense of this view is Singer (1975, chap. 1).
- ³¹ This is consistent with thinking that we are at least permitted to count the interests of beings for somewhat less in virtue of their possessing somewhat less sophisticated cognitive capacities. Of course, this sort of view requires that we accept that some humans have lesser moral status than others in virtue of possessing more limited cognitive capacities. To the extent that one finds this implication implausible, accepting the core commitments of effective altruism that I have identified generates pressure to accept that equivalent interests must always count equally, including across species lines.
- ³² This conclusion seems consistent with the way that the effective altruism movement in fact operates, and with how many effective altruists think about the issues. In addition to Animal Charity Evaluators, which focuses exclusively on organizations that work to reduce non-human animal suffering, some effective altruist organizations have animal-focused initiatives within them, along with efforts focused on other kinds of issues. The Open Philanthropy Project, for example, operates the Animal Welfare Fund (see <https://app.effectivealtruism.org/funds/animal-welfare>), in addition to funds focused on global health and development and long-term risks.
- ³³ Dylan Matthews (2015) reports that at the 2015 Effective Altruism Global conference, a number of attendees who believe that mitigating existential risks, and in particular those from artificial intelligence, should be a top effective altruist priority, described global poverty as a ‘rounding error’ in comparison with how much is at stake, morally speaking, in the effort to mitigate existential risks.
- ³⁴ For discussion, see Roberts (2011).
- ³⁵ I refer to Rawls here only to highlight what I take to be the structural similarity between the role of the overlapping consensus on core liberal principles in his view, and the role of commitment to the core claims that I have identified within effective altruism. As I go on to highlight below, noting this similarity is helpful because it allows me to highlight a contrast that I claim makes my account of the structure of the effective altruism movement especially appealing.

References

- Berg, Amy. 2018. "Effective Altruism: How Big Should the Tent Be?" *Public Affairs Quarterly* 32: 269-87.
- Berkey, Brian. "The Demandingness of Morality: Toward a Reflective Equilibrium." *Philosophical Studies* 173: 3015-35.
- _____. 2018. "The Institutional Critique of Effective Altruism." *Utilitas* 30: 143-71.
- _____. 2019a. "Collective Obligations and the Institutional Critique of Effective Altruism: A Reply to Alexander Dietz." *Utilitas* 31: 326-33.
- _____. 2019b. "Collective Obligations and Demandingness Complaints." *Moral Philosophy and Politics* 6: 113-32.
- _____. Forthcoming. "Effectiveness and Demandingness." *Utilitas*.
- Bostrom, Nick. 2002. "Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards." *Journal of Evolution and Technology* 9.
- _____. 2013. "Existential Risk Prevention as Global Priority." *Global Policy* 4: 15-31.
- Bowen, Dillon. 2016. "The Economics of Morality." *Journal of Practical Ethics* 4: 80-103.
- Gabriel, Iason. 2017. "Effective Altruism and its Critics." *Journal of Applied Philosophy* 34: 457-73.
- Gomberg, Paul. 2002. "The Fallacy of Philanthropy." *Canadian Journal of Philosophy* 32: 29-66.
- Herzog, Lisa. 2016. "Can Effective Altruism Really Change the World?" *OpenDemocracy.net*: <https://www.opendemocracy.net/transformation/lisa-herzog/can-effective-altruism-really-change-world>.
- Kolers, Avery. 2016. "Social Movements." *Philosophy Compass* 11: 580-90.
- Lichtenberg, Judith. 2015. "Peter Singer's Extremely Altruistic Heirs." *New Republic*: <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: 269-83
- _____. 2015. *Doing Good Better: How Effective Altruism Can Help You Help Others, Do Work that Matters, and Make Smarter Choices About Giving Back*. New York: Penguin Random House.
- _____. 2019. "The Definition of Effective Altruism." In *Effective Altruism: Philosophical Issues*, 10-28. ed. H. Greaves and T. Pummer. Oxford: Oxford University Press.
- Matthews, Dylan. 2015. "I Spent a Week at Google Talking with Nerds about Charity. I Came Away... Worried." *Vox*: <https://www.vox.com/2015/8/10/9124145/effective-altruism-global-ai>.
- Parfit, Derek. 1984. *Reasons and Persons*. New York: Oxford University Press.
- Pummer, Theron. 2016. "Whether and Where to Give." *Philosophy and Public Affairs* 44: 77-95.
- Pummer, Theron and MacAskill, Will. Forthcoming. "Effective Altruism." *International Encyclopedia of Ethics*.
- Rawls, John. 1993. *Political Liberalism*. New York: Columbia University Press.
- Roberts, Melinda A. 2011. "An Asymmetry in the Ethics of Procreation." *Philosophy Compass* 6: 765-76.
- Sidgwick, Henry. 1907. *The Methods of Ethics*, Seventh Edition. New York: Macmillan and Co.
- Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1: 29-43.
- _____. 1975. *Animal Liberation: A New Ethics for our Treatment of Animals*. New York: HarperCollins.
- _____. 2015. *The Most Good You Can Do: How Effective Altruism is Changing Ideas About Living Ethically*. New Haven: Yale University Press.
- Skelton, Anthony. 2016. "The Ethical Principles of Effective Altruism." *Journal of Global Ethics* 12: 137-146.
- Srinivasan, Amia. 2015. "Stop the Robot Apocalypse." *London Review of Books* 37: 3-6.
- Thomson, Judith J. 1985. "The Trolley Problem." *Yale Law Journal* 94: 1395-1415.
- Unger, Peter. 1996. *Living High and Letting Die: Our Illusion of Innocence*. New York: Oxford University Press.
- Varner, Gary E. 2012. *Personhood, Ethics, and Animal Cognition: Situating Animals in Hare's Two-*

Level Utilitarianism. New York: Oxford University Press.
Von Platz, Jeppe. 2016. "Robin Hood Justice: Why Robin Hood Took from the Rich and Gave to the Poor (and We Should Too)." *Public Affairs Quarterly* 30: 149-69.