

Effective Altruism: How Big Should the Tent Be?

Amy Berg – Rhode Island College
forthcoming in *Public Affairs Quarterly*, 2018

The effective altruism movement (EA) is one of the most influential philosophically savvy movements to emerge in recent years. Effective Altruism has historically been dedicated to finding out what charitable giving is the most overall-effective, that is, the most effective at promoting or maximizing the impartial good. But some members of EA want the movement to be more inclusive, allowing its members to give in the way that most effectively promotes their values, even when doing so isn't overall-effective. When we examine what it means to give according to one's values, I argue, we will see that this is both inconsistent with what EA is like now and inconsistent with its central philosophical commitment to an objective standard that can be used to critically analyze one's giving. While EA is not merely synonymous with act utilitarianism, it cannot be much more inclusive than it is right now.

Introduction

The effective altruism movement (EA) is one of the most influential philosophically savvy movements to emerge in recent years. Members of the movement, who often refer to themselves as EAs, have done a great deal to get philosophers and the general public to think about how much they can do to help others. But EA is at something of a crossroads. As it grows beyond its beginnings at Oxford and in the work of Peter Singer, its members must decide what kind of movement they want to have. In particular, they must decide what it truly means to be an *effective* altruist. Effective Altruism started out with what I will call a *little-tent* approach to effectiveness. Here, for something to count as “effective,” it must be *overall-effective*; that is, it must be effective *at promoting, securing, or maximizing the impartial good*. This is the sense of EA that has led people to see it as coextensive with utilitarianism.¹ (I’m leaving open for now whether little-tent EA requires us to maximize the impartial good or merely promote it—but we’ll return to that question toward the end of this paper.) The other possible approach is the *big-tent* approach. Here, “effective” merely means *value-effective*, that is, effective at *promoting, securing, or maximizing some value*. “Value” is more broadly construed than “impartial good”: people value art, political power, and community, even when they are known to contribute very little to the impartial good.

So far, EA has generally confined itself to a little tent. For example, in defining EA, Wiblin excludes from the movement those who believe that we ought to do good but who don't think this obligates us to try to maximize the impartial good.² By excluding these people, Wiblin effectively defines EA as a movement dedicated to overall-effectiveness.³ The document "What Is Effective Altruism?," co-written by a range of EAs, excludes value-effectiveness even more explicitly: "Effective altruists focus on improving lives. . . . There might be other things of value as well--promoting art, or preserving the natural environment--but that's not what effective altruism is about."⁴

But there are signs that some EAs would like to change the size of their tent. Gabriel, for example, recommends that EA broaden the kind of advice it gives, permitting "advisees to choose the values they consider morally relevant (within certain parameters) and produce recommendations that optimize the impact" on these values.⁵ This version of EA appears to take value-effectiveness as its criterion.⁶ Gabriel moves to the big tent because he believes that this version of EA is not susceptible to what he sees as the weaknesses of the little tent.⁷ Conversations I've had with EAs indicate that Gabriel isn't alone here; some people who are uncomfortable with EA's roots in utilitarianism are still attracted to the movement for its focus on evidence-based giving. They think of themselves as EAs, and yet they want to give to causes because they align with their values.

In the debate about which criterion of effectiveness to adopt, one side has a clear advantage. Broadening EA to include value-effectiveness would require abandoning some of EA's central philosophical commitments: most notably, the view that we should measure our giving using an objective standard. While this does *not* mean that EA membership is limited to consequentialists, it does mean that EA cannot include all who are interested in doing good, and it may limit the public appeal of EA. For EA to be EA, the tent must stay little.

1. Tents and Values

To see why value-effectiveness is incompatible with EA, consider:

Reproductive Access I: As a woman, Anne values reproductive rights. The value Anne wants to promote is securing reproductive rights for everyone. Therefore, Anne supports charities that are effective at ensuring reproductive access for very poor women in the developing world.

Reproductive Access I is clearly compatible with big-tent EA. Anne has a value that's important to her—reproductive rights—and she's done her research as to what can most effectively promote that value. She's thought about why she holds that value: she cares about the good, and she cares about increasing the happiness, status, and agency of women. But it's a value Anne holds because of her own life experiences, her projects, and so on, not because she believes it's the best way to pursue overall-effectiveness.

So is *Reproductive Access I* compatible with little-tent EA? Here, Anne is pursuing the value she cares about in a way that does quite a lot to promote the impartial good, since her giving targets those in great need. And, depending on the organizations Anne directs her giving to, her giving could be compatible with the recommendations of at least one EA group. The Life You Can Save (TLYCS) includes two charities dealing primarily with aspects of reproductive care (the Fistula Foundation and Population Services International) on its list of recommended charities.⁸ In order for Anne’s giving to be overall-effective, she should choose to direct her giving to one or both of those charities. If she does, then so far, her giving in *Reproductive Access I* is the kind that is encouraged by little-tent EA.

But on the other hand, neither GiveWell nor Giving What We Can (GWWC) includes a charity dealing primarily with reproductive care on its list of top charities (nor even among the runners-up).⁹ These prominent EA organizations disagree with TLYCS about whether any charities focusing on reproductive access are among the *greatest* contributors to the impartial good. If *these* organizations, not TLYCS, are correctly interpreting the available evidence, then there are other, more overall-effective ways that Anne could spend her money than by giving to reproductive-rights charities. Plus, even if Anne’s *giving* is itself effective, her *motivation* may be to give according to her values, not according to what is most overall-effective. Anne’s giving may live up to EA standards, but Anne herself might not. Thus even *Reproductive Access I*—even giving that helps save very poor women from suffering and death—cannot receive the full-throated support of all EAs.

Now consider:

Reproductive Access II: As an American woman, Beatriz values reproductive rights, and she’s disturbed by the ways in which reproductive access has been curtailed by recent laws and policies. The value she wants to promote is securing reproductive rights for Americans. Therefore, Beatriz supports charities that are effective at providing reproductive access for women in the United States and that mobilize to oppose the curtailment of reproductive access.

If *Reproductive Access I* was probably incompatible with little-tent EA, *Reproductive Access II* is definitely incompatible with it. While in *Reproductive Access I*, Anne’s donation went to help the very poor, very few American women qualify as “very poor.” Giving to support reproductive care in the United States will do much less to promote the impartial good than giving to reproductive care more broadly.

And yet, like *Reproductive Access I*, *Reproductive Access II* is compatible with big-tent EA. Big-tent EA licenses giving that supports one’s values. While “promoting reproductive access in the United States” is more geographically specific than “promoting reproductive access generally,” it’s not ad hoc. Beatriz came to hold that value for the same reasons as Anne came to hold her value in *Reproductive Access I*: it’s based on her identity, her projects, and interests. If

Beatriz's value seems less legitimate, that might be because it does less to promote the impartial good—but obviously that kind of requirement is ruled out in the big tent. If *Reproductive Access I* is permissible for big-tent EA, so is *Reproductive Access II*.

And so is:

Reproductive Access III: Chuck lives in a wealthy community in which health care is generally good. Still, some women receive care that is merely adequate rather than absolutely top-notch. Chuck's little sister works for an organization that works toward improving the quality of reproductive care in their community. The values Chuck wants to promote are reproductive access and support for the projects of people he knows and cares about. Therefore, Chuck supports his sister's organization, which is effective at securing reproductive rights within their community but which limits its focus to their community specifically.

We are now far removed from where we started. *Reproductive Access III* is clearly not the kind of giving licensed by little-tent EA. Chuck is giving to women who are already extremely well-off by global standards, and who in fact already have some access to reproductive care.

And yet big-tent EA should support *Reproductive Access III*. Chuck is, once again, operating on his values: the value of supporting his associates and the value of securing reproductive access. There is no difference between the structure of this case and the structure of the two previous cases. Although the giving in *Reproductive Access III* is extremely inefficient by little-tent standards, big-tent EA has no reason to criticize giving in the style of *Reproductive Access III*. If EAs think that the giving in *Reproductive Access I* and/or *II* seems fine (maybe not optimal, but at least reasonably good), but the giving in *Reproductive Access III* is off the table, then they must explain: Why was giving according to one's values, rather than the impartial good, permissible in Anne's and/or Beatriz's case, but not in Chuck's? What is it about Chuck's values that's out of line? But it's not clear what resources big-tent EA has to draw that line, since it cannot decide according to whether Chuck's values promote the impartial good.

The moral of these three stories is clear: once we start down the path of allowing giving according to one's values, giving quickly comes to look nothing like the giving EAs currently promote. Because of the wide range—in content and specificity—of values a person can reasonably have, a value-effective EA would encompass a range of giving, from that which targets the very poor to that which stays within a wealthy community.

What, then, is big-tent EA against? Consider:

Reproductive Access IV: Dahlia is a mean person who wants those around her to suffer. She's concluded that one way to make people suffer is to make sure they can't afford reproductive care. Dahlia gives extensively to

organizations specifically intended to prevent access to reproductive care.

In other words, big-tent EA can surely condemn giving according to nefarious values—giving intended to cause harm, suffering, or pain. Dahlia may be “effective,” but she’s no altruist. The values of *Reproductive Access IV* are not values directed at any reasonable conception of the good, impartial or not.

Is Chuck like Dahlia? In *Reproductive Access III*, he’s failing to prevent harms since he’s taking money he could have spent on cures for schistosomiasis and spending it on relatively well-funded causes instead. As long as Chuck knows that his money could be doing more, he may be demonstrating a lack of regard for those in need. But if that’s true of Chuck, then it’s also true of Beatriz, who is demonstrating a lack of regard for those in need outside the United States. And it’s even true of Anne, who is demonstrating a lack of regard for the very poor in greatest need (for whom access to reproductive care may not be the first priority). If a lack of regard for those who are worse-off is your complaint against Chuck, then it’s hard to see how you aren’t just back in the little tent.

What separates Chuck (and Anne and Beatriz) from Dahlia must be that his values are not themselves nefarious. If the world’s more pressing problems were taken care of, little-tent EA would certainly allow us to aid members of wealthy communities. Improving reproductive access in wealthy communities certainly contributes to the impartial good; it just doesn’t contribute as much as other options would, given the situation today. Clearly, there’s nothing in these values that’s *inherently* bad; they’re just at odds with little-tent EA given the context we’re currently in. Even big-tent EA ought to be able to condemn truly nefarious values (using something like Gabriel’s “parameters” for determining which values are morally acceptable), but the values at stake in *Reproductive Access III* aren’t nefarious.

Besides the pursuit of nefarious values, consider cases such as:

Reproductive Access V: As an American woman, Evelyn places some value on reproductive rights, and she’s generally against the recent attempts to curtail them, but it’s not her first priority. On her commute home, Evelyn hears a particularly upsetting story about reproductive rights on the news. When she gets home and checks the mail, she sees that a reproductive-rights organization has sent her a flyer requesting donations. Because she’s upset, she writes them a check without bothering to research whether they’re more effective than other reproductive-rights organizations at mobilizing against the curtailment of these rights.

Evelyn’s giving here is not value-effective, given her values. First, the value of reproductive rights isn’t especially important to her; she’s upset, and she’s acting on a spur-of-the-moment emotional response. Second, she’s not trying to pursue that value effectively. She doesn’t know, nor has she tried to find out, whether the charity she’s giving to is a good way to support her goals. In *Reproductive Access V*,

Evelyn is approaching instrumental irrationality, since she's pursuing a goal with no idea whether the means she's chosen will get her there.¹⁰

2. Problems with Big-Tent EA (from an EA Standpoint)

So *Reproductive Access IV* and *V* are the kinds of cases big-tent EA condemns—cases in which agents badly pursue values, or pursue bad values, or pursue values they don't care about. But big-tent EA cannot condemn any of *Reproductive Access I-III*. Because the actors in those cases each acted according to what would most effectively accomplish their values, their actions should be praised by big-tent EA.

But this shows that big-tent EA is not really EA at all. Big-tent EA is not compatible with *what* EA is; but it is also not compatible with *why* EA is that way. A big tent is neither compatible with the positions taken by major EA groups, nor with the reasons given for those positions. Effective Altruism may not be as little a tent as some have supposed, but it can only preserve its distinctive philosophical commitments—most notably, its commitment to an objective standard for giving—if it stays a little tent (although, as we will see, not as little as some might think).

2.1 What EA Is

Big-tent EA is not compatible with EA as it currently exists. Although Gabriel and others have argued that EA should help donors to choose their own values and maximize their giving in pursuit of those values, the major EA organizations do not count value-effective giving as effective giving. Giving What We Can's pledge, for example, requires its members to "give at least ten percent of what [they] earn to whichever organizations can most effectively use it to improve the lives of others."¹¹ Giving to combat climate change is generally ruled out, on the grounds that "it is very expensive to make any headway" on this issue.¹² There is no exception made if fighting climate change is an important value to you. Peter Singer's group The Life You Can Save (TLYCS) runs a similar pledge, and its guidelines are even clearer about what doesn't count: you cannot use money donated in your own community, even if it helps the poor, to fulfill your pledge to donate, since the poor near you are almost certainly not those in greatest need.¹³ Once again, what matters is overall-effectiveness. Finally, the Centre for Effective Altruism (CEA) criticizes David Geffen's donations to the Lincoln Center and the Museum of Modern Art, pointing out that "much more good could have been accomplished if that money had been spent on more pressing needs."¹⁴ If the CEA cared about value-effectiveness rather than overall-effectiveness, Geffen's gifts (assuming they represent his values) would be praiseworthy, rather than criticizable.

2.2 Why EA Is That Way

So big-tent EA is generally incompatible with the guidelines for giving espoused by

major organizations in the movement. But maybe that's not a problem. Movements evolve. "Earning to give" was once a major part of the EA organization 80,000 Hours; now it's downplayed.¹⁵ As more people become interested in EA and add their perspectives, talents, and views to the movement, EA will certainly make some further changes.

But movements can only change so much while still retaining the core philosophical commitments that make them coherent as movements. In considering whether and how EA should grow, Cotton-Barratt notes that a movement with more popular ideas will have a different growth pattern from that of a less appealing movement.¹⁶ Excising EA's focus on overall-effectiveness would, almost by definition, make it a more popular movement. Those who think overall-effectiveness is the right criterion on which to give would still be free to do so; those who think value-effectiveness is the right criterion would be newly welcome. But Cotton-Barratt warns that this tool should be used carefully: "The value of spreading a different movement may be substantially different, and it could be too easy to throw the baby out with the bathwater by removing unpopular components."¹⁷ So what's EA's "baby"? We need to know more about the distinctive philosophical commitments of EA.¹⁸

Effective Altruism's crucial philosophical commitment is that it provides a single objective standard by which we can analyze our giving. On its own, this objective standard could be compatible with either the big or the little tent. The little tent has an objective standard: Does your gift do the most it can (according to your knowledge at the time) to promote the impartial good? The big tent has an objective standard too: Does your gift do the most it can to promote *your values*, whatever those happen to be? Both of these standards are objective in some sense. Little-tent EAs have developed lists that purport to give an objective answer about how you can give most overall-effectively (although they acknowledge that these lists are based on incomplete data). While it may be more difficult to determine what your values really are, and perhaps more difficult to figure out how to most effectively promote them, there are objectively true and false answers in the big tent too. (In *Reproductive Access V*, Evelyn's spur-of-the-moment donation was objectively a bad way to promote her values.) But the standard of promoting the impartial good—the standard connected with overall-effectiveness—best captures the other philosophical commitments of EA.

First, this standard is more fully objective. Value-effective giving is assessed by an objective standard of how well you succeed at promoting a *subjectively held* value. Since value-effective EA would not be in the business of criticizing values (unless the values were truly nefarious), it would not hold all its members to precisely the same standard--which overall-effective EA does, since its standard does not depend on an individual's values.

Second, holding everyone to the same standard allows for your values to come under critical scrutiny from your peers. Wiblin warns that if EA's emphasis on justifying your plans "to a high standard" were diluted, by allowing "any plan to improve the world that isn't really ineffective," then "people [in EA] would stop

feeling any pressure from their peers to . . . change approaches.”¹⁹ If EA shifts to a bigger tent, there is no longer reason to pressure peers to give according to anything but their own values.

Third, EAs’ commitment to the impartial good allows them to avoid various psychological phenomena they see as problematic. Those who give overall-effectively are the kind of people who will, according to Hutchinson, display “open-mindedness” since they will be willing to give to a wide range of charities.²⁰ Relatedly, overall-effective giving is supposed to avoid bias. The Effective Altruism Foundation claims that if we become emotionally attached to a particular cause, “it might compromise our ability to evaluate causes objectively.”²¹ Since there is so much need in the world, and our judgments about how to best meet that need are still uncertain, we should give according to a fully objective standard. This is how we can avoid confirmation bias—if a project, or strategy, becomes particularly important to me, then I might unintentionally overestimate how effective that project or strategy is.

Finally, those who are committed to overall-effectiveness will (theoretically) do good in perpetuity.²² Suppose that Congress passes a law that guarantees full funding for reproductive care for all Americans. Beatriz (from *Reproductive Access II*) no longer has to worry about the value that drove her giving. She might stop giving altogether. Someone who is committed to overall-effectiveness, on the other hand, will switch causes once the top cause has received enough funding or is otherwise no longer an issue. Of course, this kind of concern only makes a difference at the margins. Many of the causes that people donate to appear destined to remain problems indefinitely, and many values people hold don’t have an upper limit (is there such a thing as completely fulfilling the need for art?). But the more granular a value is, the more likely that the work I can do to promote that value will eventually be exhausted (there’s only so much Chuck can do to support his sister in *Reproductive Access III*).

My purpose here is not to determine whether these philosophical commitments are the *right* ones. The Effective Altruism Foundation claims that our biases get in the way of overall-effective giving; EA’s critics might retort that those “biases” are guides to discovering your own values and projects, to thinking critically about how you want to live your life. Wiblin values a single, objective standard for giving; EA’s critics might hold that that standard narrowly constrains the values we ought to be allowed to promote.²³ But it’s clear from this discussion of EA values that EA would lose a lot of what makes it EA if it were to shift from overall-effectiveness to value-effectiveness.

We can see this even more clearly when we consider what a big-tent EA would really look like. For one thing, big-tent EA would have a marked change in focus. It would still want to encourage everyone to give, but this encouragement would be centered around individual values, not the impartial good. So EA would encourage those who don’t give at all to start giving according to their own values. Effective Altruism could look into why some people give while others don’t and could conduct persuasion campaigns to get more people to give. (That big-tent EA

uses a criterion of value-effectiveness might make it more attractive STET since there would be no requirement to restrict giving in ways that some people find alienating.) Rather than avoiding what little-tent EA sees as cognitive biases, big-tent EA might play off of those so-called biases (seeing them instead as projects, commitments, aspects of identity, and so on).

Second, big-tent EA organizations would focus on helping potential donors to understand what kind of giving would best support their values, rather than on encouraging people to give to the few overall-effective charities. For someone whose values are like those of Beatriz in *Reproductive Rights II*, EA organizations could provide a list of reproductive-rights charities doing effective work in the United States. For the David Geffens of the world, EA would need to research effective arts organizations in New York. And so on. There might be peer pressure among EAs to give *somewhere*, but there would not be peer pressure to give anywhere in particular or to revise one's values.

Finally, in service of these goals, big-tent EA could help people better understand what their own values are. Chuck might be faced with a choice in *Reproductive Access III*: Should he give to the reproductive-rights organization his sister works for, or should he give to a more effective organization serving his community? Big-tent EA could help him make the choice between these organizations by running surveys, administering quizzes, or holding counseling sessions that would lead him to understand what his own values are and how important each value is to him.

Seen this way, it becomes clear that big-tent EA is simply a sort of cheerleader for encouraging people to give however they see fit, as long as it's instrumentally rational. If EA evolves like this, it will no longer be a distinct movement. Instead, it will stand for a pretty commonsense belief: that if two charities are directed at the exact same (granularly specified) value, you should give to the more effective one. Big-tent EA is indistinct from commonsense beliefs about how one ought to promote one's values.

In fact, this kind of EA might be worse than useless. Since big-tent EA would no longer stand for a distinctive philosophical belief, and since it never focused on a particular cause, it would be in danger of becoming otiose. The Red Cross may not stand for a philosophically controversial position, but at least it's staked out a cause area. Big-tent EA would be more like a clearinghouse for lots of different causes and lots of different values. If the costs of running EA organizations (paying staff, maintaining web servers, creating and disseminating information) outweigh the benefits of having such a clearinghouse, we might just want to cut out the middleman entirely.

There are three things EAs should take away from this discussion. First, the cases of *Reproductive Access I-III* show that if EA admits value-effective giving, then many different ways to give will count as effective, even if that giving exclusively benefits the well-off. Cases that EAs frequently criticize, such as Geffen's gifts to the arts, look laudable from a big-tent perspective. Second, if we look at the statements of flagship EA organizations and people, it's clear that value-effective giving is not

what they have in mind. But, of course, that's exactly what Gabriel and others advocate changing. So, third, EAs should be aware of how many of their philosophical commitments they would give up by changing to a big-tent EA. Of course, EAs own their movement, and it's ultimately up to them to decide what kind of movement they want, but value-effectiveness represents a substantial break with the philosophical commitments that EAs presently hold. They cannot both preserve these commitments and enlarge their tent.

3. A Big Tent for Little-Rent Reasons?

Here's one other possibility. The founders of EA are, for the most part, fairly standard act utilitarians. For this group, there's nothing wrong with being quiet about your real motives, or maybe even actively engaging in deception, *if* doing so maximizes the impartial good. This could lead to an EA movement that is *both* big- and little-tent: its public face encourages people to give more in whatever way is value-effective, while its private face remains focused on the overall-effectiveness that was EA's original target. If big-tent EA could get people excited about giving, over time, it could expose them to "real" EA—to the idea that the giving they're doing could be much more overall-effective.

This might lead to some rhetorical changes in how EA approaches charitable giving. Remember when the "Ice Bucket Challenge" was popular back in 2014? People were dumping buckets of water on themselves in order to raise money to fund research into ALS (amyotrophic lateral sclerosis; Lou Gehrig's disease). A number of EAs vocally criticized the Ice Bucket Challenge, arguing that the funds it was raising for ALS research were "cannibalizing" funds that could have gone to more effective charities.²⁴ If EAs were to adopt a big-tent strategy for little-tent reasons, they might approach the Ice Bucket Challenge differently, either by co-opting the ice bucket strategy for overall-effective charities, or by trying to recruit people who had already done the Ice Bucket Challenge for ALS into becoming donors to overall-effective charities afterward. In fact, MacAskill implies that EAs could co-opt this strategy, in an article entitled "This Week, Let's Dump a Few Ice Buckets to Wipe Out Malaria Too" that was intended as a follow-up to his original piece criticizing the Ice Bucket Challenge; other effective altruists made similar arguments.²⁵ The Ice Bucket Challenge got people temporarily excited about giving; EA could pick up on that enthusiasm, redirect it, and make it more permanent by having a big tent in front and a smaller tent in back.

Insofar as this strategy involves dishonesty, it's bound for trouble. It may be morally acceptable to the strict act utilitarians who were the first EAs, since this kind of duplicity wouldn't be impermissible (in fact, would be obligatory), on their view, as long as it maximized overall utility. But it's not clear that even a strict utilitarian calculus would come out in favor of a big tent for little-tent reasons. People tend not to like feeling like they're being manipulated. If people came to believe that EA had been rhetorically duplicitous—if they felt used, manipulated, or tricked by a little-tent organization with a big-tent face, that might turn them

away from EA. Even worse, EA might be a cautionary tale used by those who are looking for an excuse to avoid all charitable giving whatsoever (“See, I told you those do-gooders couldn’t be trusted!”). This strategy might also diminish EA’s appeal for non-act-utilitarian members, since at least some more deontologically inclined little-tent EAs might see something wrong with using big-tent EA supporters as mere means to an eventual little-tent EA end. So if there’s room in a little tent for people who are not strict act utilitarians to be EA (and as I’ll lay out below, I think there is), those people may be uncomfortable with a movement that is dishonest about its true theoretical beliefs.

But a big tent out front, with a little tent in the back, need not involve dishonesty. A different version of this strategy would counsel EAs to be open and honest about their beliefs. EAs could say: “We think that everyone should give according to overall effectiveness, but we recognize not everyone agrees with us. We still think you should give, even if you do so less effectively.” This one is risky too. First, I think it runs the risk of being condescending. If EAs are encouraging others to give how they think best, but at the same time are saying that other views on giving are misguided, that could make non-EA givers feel diminished. Even if EAs can successfully not condescend to the people they’re trying to attract, I have doubts about the effectiveness of this strategy. I don’t know that it’s very motivating to hear that you should give, but that the way you want to give is less valuable or praiseworthy than the way EAs give is. The more honest EAs are about how they think of effectiveness and why it matters, the harder it will be for them to attract people who are open to giving but inclined to disagree about how to give.

So I have my doubts as to whether a big tent for little-tent reasons could be a successful strategy for EAs to pursue. But whether EAs pursue this strategy openly or dishonestly, it’s just a strategy. If we’re investigating the possibility of a big tent for little-tent reasons, then the *theoretical* argument has already been won by little-tent EA; we’re just arguing now about which *pragmatic* strategy is more effective. All of the philosophical work is being done by the little-tenters, even if EA appears to have changed. Even if EA were to adopt this strategy, which I think is dicey both morally (if it involves dishonesty) and pragmatically (if it can’t motivate), the tent that matters would stay the same size.

4. What about Different EA-Endorsed Charities?

I’ve argued that EA cannot evolve beyond its current view of effectiveness as overall-effectiveness. But is this really consistent with EA practice so far? Perhaps EAs really do permit giving in a value-effective way. Evidence for this might be the split between EA organizations focused on animal causes (such as Animal Charity Evaluators, or ACE) and those whose recommended charities mostly help humans (such as GiveWell and GWWC). The human-focused charities disagree with each other too. As we saw back in *Reproductive Access I*, TLYCS endorses the Fistula Foundation, but GWWC and GiveWell do not.

We shouldn’t overestimate this split, though: these organizations have a lot

of members in common, share many guiding values, and generally support rather than rival each other. But they do have different areas of focus, and someone donating 10 percent of his income to effective charities has to decide whether to give that money to animal- or human-related causes. If it's permissible for an EA to give to either human or animal charities, then it might be permissible because value-effectiveness is the right criterion for deciding where to give. Conversely, if overall-effectiveness is the right criterion of effectiveness, then GiveWell and ACE can't both be right. Either all EAs should be working exclusively on human causes, or they should all be giving exclusively to animal charities.

But value-effectiveness isn't what's doing the work here. Animal Charity Evaluators pursues animal-related causes because *we don't know* whether animal or human charities are the most overall-effective ways to give. Animal Charity Evaluators, GiveWell, and GWWC share a commitment to overall-effectiveness, but those who give to ACE-recommended charities believe that these may be the most overall-effective. Animal Charity Evaluators and GiveWell, and also GiveWell and GWWC, disagree not about whether overall- or value-effectiveness is the right sense of effectiveness; they just may have disagreements about how to calculate what is most overall-effective. This is, very reasonably, due to the extreme difficulty of calculating what does the best job of maximizing happiness and relieving suffering, particularly where animals are concerned. But if the research comes in, and it turns out that animal charities are clearly more effective, we would expect EA donors to human causes to switch to animal causes.

5. What about a Medium Tent?

But what if that doesn't happen? What if the evidence comes in, and animal charities are more effective at relieving suffering than human-focused charities are, but donors to GWWC-endorsed charities *don't* switch their giving? Donors to ACE and GWWC may be giving in order to promote the impartial good and yet disagree about what that impartial good is, even after all the evidence is in. If they do, this suggests that there are more tents to try out.

Hutchinson proposes a kind of medium tent in a discussion of why EA is "cause-neutral" (by which she means that EAs commit to giving according to the impartial good rather than on the basis of "personal connections" to a cause).²⁶ She writes:

But being cause neutral is not the same as valuing everything. Different people value different things, and there are numerous plausible ethical systems. You might or might not value complexity and diversity; you might or might not value non-human animals; you might or might not value lives which haven't yet to come into existence [*sic*]. The more inclusive your moral system, the more likely you are to feel that moral systems that exclude some of the things you value are partisan. For example someone who values non-human animals might feel that someone whose moral

system does not attribute value to non-human animals was being partisan, just as someone who helps those near to them rather than those far away is being partisan.²⁷

This helps us to see what a medium tent might look like.²⁸ Like the big tent, the medium tent takes as its definition of effectiveness overall-effectiveness. So once again, the medium tent includes only EAs who are committed to giving in ways that are effective at promoting, securing, or maximizing the impartial good. But the medium tent allows for disagreement (such as that between ACE and GWWC) about what the impartial good consists of. This would not be as big as the big tent, which licenses giving according to your values, even when they diverge from what you believe would best promote the impartial good. But it's not as little as the little tent, since the medium tent includes people who give to many different kinds of causes for different reasons. The medium tent offers more theoretical consistency than the big tent and more inclusivity than the little tent. But there are a few things potential medium-tenters should keep in mind.

First, many EAs probably think some reasonable value systems are worse than others. Look at how Hutchinson describes the different value systems—some “value complexity and diversity” (usually seen as good things) while others do not. Some value non-human animals, while some are “partisan,” not exactly a term of praise in EA circles.²⁹ This is hardly a neutral weighing of pros and cons, suggesting that coexistence in the medium tent may not always be harmonious.

And that may be as it should be. I won't construct a full case series for the medium tent, but we can see how criticisms will be parallel to those leveled against the big tent in *Reproductive Access I-III*. Consider:

Famine Relief: Francis's value system holds that we have moral obligations only to humans. The charity that is the most effective at relieving human suffering in a recent famine is a charity that provides factory-farmed meat to those without food. In accordance with his value system, Francis donates to this charity, thereby increasing the suffering of some number of animals.

Francis's value system seems like the kind of system about which Hutchinson would say reasonable people could disagree. So this kind of giving seems to fit within the medium tent. But the benefit to Francis's version of the impartial good will come at a cost to the *standard EA* view of the impartial good, since Francis's giving causes additional animal suffering. This kind of medium-tent giving may actually be *worse*, by little-tent standards, than Anne's giving in *Reproductive Access I*. At least there, Anne was doing nearly as much good as she could.

In looking at the series of big-tent cases, we saw that it was going to be very difficult to draw a line between value-effective giving that counted as EA and value-effective giving that did not. The same is true here. There are plenty of reasonable value systems that have a picture of what counts as the “impartial good” that is different from the little-tent EA picture. Once again, either EA must make room for

value systems that are sharply at odds with what would have seemed to be its core theoretical commitments, or it must find some way to draw a line between reasonable value systems that can count as EA and those that can not. There may be other possible medium tents, but the problems shared by the big and medium tents suggest that parallel problems will recur for any attempt to make EA more inclusive.

Finally, the costs of expanding to a medium tent have relatively small gains. The advantage of the big tent was that it allowed anyone in who was interested in giving effectively in accordance with her values, even if she acknowledged that pursuing her values wasn't the best way to achieve the impartial good. The medium tent can't include those people since it only includes those whose giving effectively promotes the impartial good (as they define it). Anyone who wants to give effectively in accordance with her projects is still outside the medium tent.

6. How Little Is the Little Tent?

If EA has to be a little tent, what do we mean by "little"? Our discussion so far has shown that you must be committed to the pursuit of (a particular understanding of) overall-effectiveness rather than value-effectiveness. This little tent doesn't exclude nonconsequentialists. There are already nonconsequentialists within EA; for one, Ramakrishnan argues that a Kamm-style analysis of giving shows that overall-effective giving is morally required.³⁰ There are familiar nonconsequentialist moral theories that are compatible with EA; Ross's duty of beneficence, although situated within a larger deontological moral theory, is a duty to maximize the impartial good.³¹ These examples show that Lichtenberg is wrong to claim that EA just *is* consequentialism.³²

But Lichtenberg isn't *very* wrong. Any nonconsequentialist who doesn't accept the overall-effectiveness reading of EA gets kicked out of the little tent. For many nonconsequentialists, beneficence *should* be a value-effective duty. It is permissible, these nonconsequentialists argue, to believe that beneficence is inseparable from one's other values, commitments, and projects.³³ For these nonconsequentialists, and possibly for some consequentialists, too, there's no room in the little tent.

So the commitment to overall-effectiveness includes utilitarians, many other consequentialists, and some nonconsequentialists, while excluding a significant number of nonconsequentialists. There's another dimension to the tent that's relevant too. At the beginning of this paper, I described a commitment to overall-effectiveness as a commitment to *promote, secure, or maximize* the impartial good. Once we see that EAs are required to give in an overall-effective way, the size of the tent will vary further depending on whether they are also required to maximize the impartial good (by giving to the level of marginal utility, as Singer describes in "Famine, Affluence, and Morality") or merely to promote it, by giving at a lower level than that.³⁴ If I give \$5 once to the Against Malaria Foundation, am I an effective altruist?

Effective Altruism organizations currently require some level of giving, although none requires maximizing the impartial good. Giving What We Can asks its members to pledge 10 percent of their income; The Life You Can Save has an income-based sliding scale.³⁵ Still, while the commitment to overall-effectiveness falls straightforwardly out of EA's philosophical commitments, a commitment to a particular level of overall-effective giving does not. To survive as a philosophically coherent movement, EA needs a commitment to the impartial good more than it needs a commitment to maximizing the good. If EA requires its members to give at a particular level, the message is still that giving *more* would be better. But if EA requires its members to give according to their values, that doesn't send a clear message that giving according to overall-effectiveness would be better. Perhaps this shows that the commitment to overall-effectiveness, not the commitment to a certain level of giving, is the true core idea of EA.

Or perhaps overall-effectiveness and a certain level of giving are both core ideas of EA, in which case EA will have to figure out what that level is and why its members are required to give that much.³⁶ Depending on how that argument shakes out, it may be that when it comes to the level of required giving, the little tent has slightly more room.³⁷

7. Conclusion

I have argued that EA must be a little tent rather than a big tent. In order to be a distinct movement that retains its core philosophical beliefs, rather than just being a cheerleader for instrumental rationality, EA must adopt a criterion of overall- rather than value-effectiveness.

Still, none of this means that EA is *right* to adopt overall- rather than value-effectiveness. Effective Altruism may fail to recognize the importance of values other than the impartial good—the importance of associative duties, the role giving can play in forming and contributing to an individual's identity and commitments, the space for personal prerogatives, and so on. This paper does not tackle the question of whether EA is the right kind of approach to take to beneficence. But since EA has made significant contributions to the philosophical debates around giving, I think it's in all our interests to clarify what EA is and what it stands for.

Rhode Island College

Notes

A previous version of this paper was presented at the 2018 Philosophy, Politics, and Economics Society Meeting, for which I received funding from the Institute for Humane Studies. Thanks to Craig Agule, Brian Berkey, Luc Bovens, Matt Braich, Brookes Brown, Alexander Dietz, Daniel Kokotajlo, Per-Erik Milam, Dan Shahaar, Jeff Sebo, Steve Swartzler, Travis Timmerman, and the editor and an anonymous referee at *Public Affairs Quarterly* for their helpful feedback.

References

- Carter, Ben, and Keith Moore. "Icy Logic: Choosing a Charity with Your Head and Your Heart." *BBC Magazine*, September 7, 2014. www.bbc.com/news/magazine-29087531.
- Centre for Effective Altruism. "Introduction to Effective Altruism." *Effective Altruism*, June 22, 2016. www.effectivealtruism.org/articles/introduction-to-effective-altruism/.
- , "Our Mistakes." *80,000 Hours*, December 2017. 80000hours.org/about/credibility/evaluations/mistakes/.
- Cotton-Barratt, Owen. "How Valuable Is Movement Growth?" Effective Altruism Forum, May 14, 2015. http://effectivealtruism.com/ea/is/how_valuable_is_movement_growth/.
- Effective Altruism Foundation. "The Benefits of Cause-Neutrality." April 11, 2017. <https://ea-foundation.org/blog/the-benefits-of-cause-neutrality/>.
- Gabriel, Iason. "Effective Altruism and Its Critics." *Journal of Applied Philosophy* 34, no. 4 (2017): 457-73.
- GiveWell. "Top Charities." n.d. www.givewell.org/charities/top-charities.
- Giving What We Can. "Frequently Asked Questions." n.d. www.givingwhatwecan.org/about-us/frequently-asked-questions.
- , "The Pledge." n.d. www.givingwhatwecan.org/pledge/.
- , "Top Charities." n.d. www.givingwhatwecan.org/top-charities/.
- , "Why We (Still) Don't Recommend GiveDirectly." November 10, 2015. givingwhatwecan.org/post/2014/02/why-we-still-dont-recommend-givedirectly/.
- Halstead, John, Stefan Schubert, Joseph Millum, Mark Engelbert, Hayden Wilkinson, and James Snowden. "Effective Altruism: An Elucidation and a Defense." April 4, 2017. Unpublished manuscript.
- Herman, Barbara. "Mutual Aid and Respect for Persons." *Ethics* 94, no. 4 (1984): 577-602.
- Hutchinson, Michelle. "Giving What We Can Is Cause Neutral." Effective Altruism Forum, April 22, 2016. <http://effectivealtruism.com/ea/is/giving-what-we-can-is-cause-neutral/>.

altruism.com/ea/wj/giving_what_we_can_is_cause_neutral/.

Krishna, Nakul. "Add Your Own Egg." *Point*, 2016. thepointmag.com/2016/examined-life/add-your-own-egg.

Lichtenberg, Judith. "Effective Altruism: A Critique." Paper presented at the Ethics of Giving Conference, University of St. Andrews, St. Andrews, UK, May 24, 2017.

The Life You Can Save. "The Best Charities for Highly Effective Giving." 2018. <https://www.thelifeyoucansave.org/best-charities>.

----- "How Much Should I Give?" 2018. www.thelifeyoucansave.org/take-the-pledge.

----- "Peter Singer Answers Common Questions about Giving." 2018. <https://www.thelifeyoucansave.org/take-the-pledge/faq>.

MacAskill, William. "The Cold, Hard Truth about the Ice Bucket Challenge." *Quartz*, August 14, 2014. qz.com/249649/the-cold-hard-truth-about-the-ice-bucket-challenge/.

----- "This Week, Let's Dump a Few Ice Buckets to Wipe Out Malaria Too." *Quartz*, August 18, 2014. qz.com/250845/this-week-lets-dump-a-few-ice-buckets-to-wipe-out-malaria-too/.

Pummer, Theron. "Whether and Where to Give." *Philosophy and Public Affairs* 44, no. 1 (2016): 77-95.

Ramakrishnan, Ketan. "Acting Wrongly beyond the Call of Duty." Paper presented at the Ethics of Giving Conference, University of St. Andrews, St. Andrews, UK, May 25, 2017.

Ross, W. D. *The Right and the Good*. Oxford, UK: Clarendon, 1965.

Schubert, Stefan. "Understanding Cause-Neutrality." *Centre for Effective Altruism*, March 10, 2017. <https://www.centreforeffectivealtruism.org/blog/understanding-cause-neutrality/>.

----- "Why the Triviality Objection to EA Is Beside the Point." *Effective Altruism Forum*, July 20, 2015. effective-altruism.com/ea/17/why_the_triviality_objection_to_ea_is_beside_the/.

Singer, Peter. "Ethics and Intuitions." *Journal of Ethics* 9, nos. 3-4 (2005): 331-52.

- . "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1, no. 3 (1972): 229-43.
- Srinivasan, Amia. "Stop the Robot Apocalypse." *London Review of Books* 37, no. 18 (2015): 3-6.
- WhatIsEffectiveAltruism.com. "What Is Effective Altruism?" 2015. <https://whatiseffectivealtruism.com/archive/>
- Wiblin, Robert. "Disagreeing about What's Effective Isn't Disagreeing with Effective Altruism." *80,000 Hours*, July 16, 2015. 80000hours.org/2015/07/disagreeing-about-whats-effective-isnt-disagreeing-with-effective-altruism/.
- . *Movement Development*. YouTube video. 17:35 min. Posted by Effective Altruism Global. November 25, 2015. www.youtube.com/watch?v=TH4_ikhAGz0.

¹ See Lichtenberg ("Effective Altruism").

² Wiblin, "Disagreeing about What's Effective."

³ See also Schubert's ("Why the Triviality Objection") response to Wiblin. Schubert appears to accept Wiblin's version of EA's "core idea" and yet argues that "large swathes of the liberal, educated part of the population" will see EA's message as "trivial/obvious."

⁴ <https://whatiseffectivealtruism.com/archive/>.

⁵ Gabriel, "Effective Altruism," 470.

⁶ Although to know the precise size of Gabriel's tent, we still need to know more about what the parameters are.

⁷ Gabriel, "Effective Altruism," 459.

⁸ The Life You Can Save, "Best Charities."

⁹ GiveWell, "Top Charities"; Giving What We Can, "Top Charities."

¹⁰ Some utilitarians hold that utilitarianism is rationally required (see Singer, "Ethics and Intuitions," 351). If that's true, then *Reproductive Access I-III* might all be irrational. But we've now moved from evaluating the rationality of a particular means to evaluating the rationality of the ends themselves. Unlike *Reproductive Access I-III*, *Reproductive Access V* is irrational on both counts.

¹¹ Giving What We Can, "Pledge."

¹² Giving What We Can, "Frequently Asked Questions."

¹³ The Life You Can Save, "Peter Singer Answers Common Questions."

¹⁴ Centre for Effective Altruism, "Introduction to Effective Altruism."

¹⁵ See Centre for Effective Altruism ("Our Mistakes") on [80,000 Hours](#).

¹⁶ Cotton-Barratt, "How Valuable."

¹⁷ Cotton-Barratt, "How Valuable."

¹⁸ There are also pragmatic reasons for EA to be cautious about the switch to a big tent. Even if EAs are tempted by this switch, Wiblin (*Movement Development*) notes that it may be easier to change a movement's brand than it is to change that brand *back*. If enlarging the tent seems otherwise desirable, this point provides a pragmatic reason to move slowly with change, if at all. But table this pragmatic point to focus on the philosophical issues at hand.

¹⁹ Wiblin (*Movement Development*).

²⁰ See Hutchinson ("Giving What We Can Is Cause Neutral") on the Effective Altruism Forum.

²¹ Effective Altruism Foundation, "Benefits of Cause-Neutrality."

²² Schubert, "Understanding Cause-Neutrality."

²³ Wiblin, "Disagreeing about What's Effective."

²⁴ MacAskill, "Cold, Hard Truth."

²⁵ MacAskill, "This Week"; Carter and Moore, "Icy Logic."

²⁶ I have avoided the term "cause-neutrality" elsewhere, because, as Schubert ("Understanding Cause-Neutrality") points out, EAs use that term to mean different things. But Hutchinson's meaning here is helpful for understanding what a medium tent would look like.

²⁷ Hutchinson, "Giving What We Can."

²⁸ Something like this alternative was suggested to me by Brian Berkey.

²⁹ Hutchinson, "Giving What We Can."

³⁰ Ramakrishnan made this claim in a talk called "Acting Wrongly beyond the Call of Duty."

³¹ Ross, *Right and the Good*.

³² Lichtenberg made this claim in a talk called "Effective Altruism: A Critique."

³³ Herman ("Mutual Aid"), among many others, develops this kind of nonconsequentialist view of beneficence. Effective Altruism's many critics often sound similar themes; see, for example, Krishna ("Add Your Own Egg") and Srinivasan ("Stop the Robot Apocalypse").

³⁴ Singer ("Famine, Affluence," 234).

³⁵ Giving What We Can, "Pledge"; The Life You Can Save, "How Much Should I Give?"

³⁶ Dan Shahar pointed out to me that if EA allows members to only give a certain amount of their income (say, 10 percent) to overall-effective charities, that perhaps their pledge could be discharged by giving 20 percent of their income to a charity that is half as overall-effective as the most effective charity. If EA allows that kind of thing, then EAs would have a little bit more discretion about where to give. But since many value-effective charities are not even a fraction as overall-effective as EA organizations' top recommended charities, this strategy won't enlarge the tent all that much.

³⁷ Pummer ("Whether and Where to Give") has argued that even if we have options about *whether* to give, we may not always have options about *where* to give. This suggests that the commitment to overall-effectiveness (that is, making sure *where* we give is an overall-effective charity) is more important than the level of giving (that is, is more important than *whether* or not we maximize the level of giving we engage in).