

# Why Ten Percent?

AMY BERG\*

## ABSTRACT

*While effective altruists (EAs) spend a lot of time researching which ways to do good are the most effective, historically many have assumed, with relatively little argument, that the benchmark for membership in the movement is a commitment to donate 10% of your earnings. This points to an asymmetry between the two halves of effective altruism: EAs tend to have relatively restricted standards for effectiveness (where to give), but they have much looser standards for altruism (how much to give). I investigate explanations for this asymmetry. While some possible justifications may work (pending empirical support), others look flimsier. I conclude that this means EA likely is, or anyway ought to be, more demanding than some of its proponents currently claim.*

## TABLE OF CONTENTS

I. AN ASYMMETRY BETWEEN <i>EFFECTIVE</i> AND <i>ALTRUISM</i> . . . . .	656
II. EXPLANATIONS FOR THE ASYMMETRY . . . . .	657
A. <i>Arbitrariness</i> . . . . .	657
B. <i>Demandingness</i> . . . . .	658
C. <i>EA as a Project, not an Obligation</i> . . . . .	660
D. <i>Research Priorities</i> . . . . .	664
E. <i>EA as a Conditional Obligation</i> . . . . .	666
F. <i>Alienation</i> . . . . .	668
III. RESOLVING THE ASYMMETRY . . . . .	670

---

\* Assistant Professor of Philosophy, Oberlin College. I received helpful comments from other participants at the Georgetown Institute for the Study of Markets and Ethics symposium on effective altruism. I'm especially grateful to Jeff Sebo for a number of thought-provoking questions and to Joshua Stein for feedback on a previous version of this paper. © 2023, Amy Berg.

I. AN ASYMMETRY BETWEEN *EFFECTIVE* AND *ALTRUISM*

Effective altruists have long had a clear, non-arbitrary standard for *where* a person ought to donate their money (and how they ought to choose a career, etc.): one should do whatever does the most good.<sup>1</sup> This means that when EAs donate in their capacity as EAs, they give to just a few organizations, such as those at the top of GiveWell’s list of recommended charities.<sup>2</sup> While EAs disagree about exactly where to give (human vs. animal suffering, near- vs. long-term well-being), they see these as questions about how to apply the standard, not about what the standard ought to be in the first place. Some giving clearly falls short of that standard: giving to an arts program in a relatively wealthy community, for example, or even training guide dogs rather than preventing trachoma in developing countries.<sup>3</sup>

At the same time, EA lacks a clear, non-arbitrary standard for *how much* a person ought to give. EA organizations (Giving What We Can, GiveWell, the Center for Effective Altruism (CEA)) have generally settled on the figure of 10% of a person’s earnings.<sup>4</sup> The 10% standard is prominent in the philosophical literature on EA too—recent work by Pummer, Berkey, MacAskill, and Mogenson and Ord, among others, all cite 10% as the standard.<sup>5</sup> It is also likely to be familiar to many EA-curious people; it’s a tithe, like the ones found in many Jewish, Christian, and Latter-day Saints traditions.<sup>6</sup> Still, while it’s common to see

1. The nature of this “ought” (whether it’s a moral obligation or an obligation you commit yourself to as part of a voluntary project you take on) is less clear; we’ll talk about that more later. Since donating earnings is the center of EA practice, that’s what this paper will focus on; it would be interesting to consider whether career choice is vulnerable to the same asymmetry.

2. Since EAs (more or less) see giving at 10% of their earnings as the benchmark for counting as EAs, this doesn’t rule out giving to other, less effective causes after they’ve given their 10% to effective charities.

3. Singer has made the claim about guide dogs in a number of places, for example in a TED talk. Peter Singer, *The Why and How of Effective Altruism*, TED: IDEAS WORTH SPREADING (2013), [https://www.ted.com/talks/peter\\_singer\\_the\\_why\\_and\\_how\\_of\\_effective\\_altruism](https://www.ted.com/talks/peter_singer_the_why_and_how_of_effective_altruism) [<https://perma.cc/HZ52-XWTT>]. EAs have also been critical of this claim. See saulius, *Fact Checking Comparison Between Trachoma Surgeries and Guide Dogs*, EFFECTIVEALTRUISM.ORG: EFFECTIVE ALTRUISM FORUM (May 10, 2017), <https://forum.effectivealtruism.org/posts/SMRHnGXirRNpvB8LJ/fact-checking-comparison-between-trachoma-surgeries-and> [<https://perma.cc/N7BX-RVBH>].

4. Giving What We Can’s pledge is here: *Our Pledge*, GIVING WHAT WE CAN, <https://www.givingwhatwecan.org/pledge> [<https://perma.cc/89AH-2JSH>]; CEA cites that pledge here: *Consider a Giving Pledge*, EFFECTIVEALTRUISM.ORG, <https://www.effectivealtruism.org/get-involved/take-the-giving-what-wecan-pledge> [<https://perma.cc/QL29-25BX>]; GiveWell cites it here: Holden, *Organizations Promoting Generous, Effective Giving*, GIVEWELL (Jan. 15, 2018), <https://blog.givewell.org/2015/01/08/organizations-promoting-generous-effective-giving/> [<https://perma.cc/95FS-7VQK>].

5. THERON PUMMER, *THE RULES OF RESCUE* 208–09 n.7 (2023) [hereinafter PUMMER, *RESCUE*]; Brian Berkey, *Effectiveness and Demandingness*, 32 *UTILITAS* 368 (2020) [hereinafter Berkey, *Effectiveness*]; William MacAskill, *The Definition of Effective Altruism*, in *EFFECTIVE ALTRUISM: PHILOSOPHICAL ISSUES* 10–28 (Theron Pummer & Hilary Greaves eds., 2019) [hereinafter MacAskill, *Definition*]; William MacAskill et al., *Giving Isn’t Demanding*, in *THE ETHICS OF GIVING: PHILOSOPHERS’ PERSPECTIVES ON PHILANTHROPY* (Paul Woodruff ed., 2018) [hereinafter MacAskill et al., *Giving*].

6. Other religious traditions have requirements to give to charity, just not necessarily 10% requirements (*tithe* means, literally, “tenth”). *Tithe*, *OXFORD ENGLISH DICTIONARY* (3d ed. 2018).

references to 10%, there are alternatives: Giving What We Can offers a “Trial Pledge” for people who want to start out at just 1% of their earnings, and Singer’s *The Life You Can Save* asks people to pledge at an income-relative rate lower than 10%. Even some proponents of the 10% standard recognize that it is arbitrary.<sup>7</sup>

It seems weird for a movement that prides itself on rational argumentation to accept an arbitrary standard, and it sets up an asymmetry between the two halves of EA. The *effective* side (*where* to give) is evidence-based and demanding. The *altruism* side (*how much* to give) isn’t often argued for and is relatively lax. Can this asymmetry be justified?

## II. EXPLANATIONS FOR THE ASYMMETRY

We’ll focus on six different explanations: there’s a non-arbitrary standard for effectiveness but not altruism; for pragmatic reasons, EA shouldn’t demand too much; EA is best seen as a discretionary project and not a moral obligation; EA should focus on research rather than donations; EA is only a conditional obligation; and a symmetrical EA would be too alienating. None of these explanations, we will see, fully justify this asymmetry. Instead, EA that asymmetrically emphasizes strict standards for effectiveness and laxer standards for altruism is deeply in tension with the movement’s core claims.

### A. Arbitrariness

One pretty straightforward explanation for the asymmetry is to look to where EA can (and cannot) make non-arbitrary claims. EAs might say that there is a non-arbitrary place to set the threshold for effectiveness: your giving should be as effective as you can make it. You’re either giving to the most effective charity, or you aren’t. There is, however, no non-arbitrary place to set the threshold of altruism. If you give 100% of your earnings to effective charities, you’re certainly an EA; if you give 1% away, you aren’t; and there’s a lot of room in between. We might as well go with 10%—clear, simple, familiar, doable. The asymmetry exists because one side is non-arbitrary and the other can’t be.

In practice, things might not be so clean on the effectiveness side of the ledger. GiveWell currently lists four top charities; Animal Charity Evaluators lists three

---

7. In a recent interview, MacAskill comments:

I don’t think there’s a privileged stopping point, philosophically. At least not until you are at the point where you are really doing almost everything you can. So with Giving What We Can, for example, we chose 10 percent as a target for what portion of people’s income they could give away. In a sense it’s a totally arbitrary number. Why not 9 percent or 11 percent? It does have the benefit of 10 percent being a round number. And it also is the right level, I think, where if you get people to give 1 percent, they’re probably giving that amount anyway. Whereas 10 percent, I think, is achievable yet at the same time really is a difference compared to what they otherwise would have been doing.

Jacob Stern, *Effective Altruism’s Philosopher King Just Wants to be Practical*, ATL MONTHLY (Sept. 29, 2022), <https://www.theatlantic.com/health/archive/2022/09/oxford-philosophy-professor-william-macaskill-effective-altruism-interview/671597/> [https://perma.cc/V2DG-TCVR].

(and 13 “standout” charities that aren’t quite as good); Giving What We Can’s pledge lacks a specific requirement beyond giving to “whichever organizations can most effectively use” the money; it isn’t clear what longtermists (EAs who focus on long-term causes, such as preventing civilizational collapse or nuclear catastrophe) are supposed to give to. In theory, EAs ought to give to the most effective organization out there, since doing anything less than that would be “absurd”; in practice, EAs give themselves considerably more wiggle room.<sup>8</sup>

Effective altruists might not see this as a problem. After all, they point out, people have reasonable disagreements about morality, and—more crucially for EAs—they disagree about the empirical facts too.<sup>9</sup> It is hard to pinpoint how much human suffering each dollar sent to the Malaria Consortium relieves, and how that compares to the Against Malaria Foundation—and things get more complicated when we try to compare current human suffering to the suffering of animals or of beings in the far future. In these conditions of uncertainty, an EA might say, of course there’s room for reasonable disagreement.

At the same time, it’s also not quite right to say that there’s only an arbitrary place to set the threshold for how much to give. A plurality, even maybe a majority, of EAs claim to be utilitarians, and so many if not most EAs accept a non-arbitrary obligation of beneficence:<sup>10</sup> something along the lines of a requirement to give to the level of marginal utility (which for most EAs is well above the 10% threshold).<sup>11</sup> EAs who are utilitarians have a clear and non-arbitrary criterion for how much to give, even as a lack of data makes it harder to tell which causes are truly the most effective. In other words, the side of the asymmetry that seemed clear is murkier than it looked, and the murkier side is a little bit clearer. This doesn’t tell us there’s no asymmetry between the two halves—just that exactly how different they are is more complicated than we might have thought.

### B. Demandingness

The many EAs who are utilitarians are likely to be familiar with demandingness objections to this moral theory. If you’re concerned that people view your moral commitments as being too demanding, and you want to start a social movement, it’s probably a good idea to make sure that movement isn’t seen as

8. WILLIAM MACASKILL, *DOING GOOD BETTER: EFFECTIVE ALTRUISM AND HOW YOU CAN MAKE A DIFFERENCE* 50 (2016).

9. Michelle Hutchinson, *Giving What We Can Is Cause Neutral*, EFFECTIVE ALTRUISM F., (April 22, 2016) [http://effective-altruism.com/ea/wj/giving\\_what\\_we\\_can\\_is\\_cause\\_neutral/](http://effective-altruism.com/ea/wj/giving_what_we_can_is_cause_neutral/) [<https://perma.cc/6D3U-SCRR>].

10. See Chris Cundy, *The 2015 Survey of Effective Altruists: Results and Analysis*, EFFECTIVE ALTRUISM F., (July 29, 2016) <https://forum.effectivealtruism.org/posts/EqgF5amS96TahanzK/the-2015-survey-of-effective-altruists-results-and-analysis> [<https://perma.cc/J9ML-VDNE>]. See also Ellen McGeoch & Peter Hurford, *EA Survey 2017*, EFFECTIVE ALTRUISM F. (2017), <https://forum.effectivealtruism.org/posts/TuzkYANf7tmQDF2Hk/ea-survey-2017-series-distribution-and-analysis-methodology> [<https://perma.cc/PB2W-GPV7>]; Brian Berkey, *The Philosophical Core of Effective Altruism*, 51 J. SOC. PHIL. 93, 98 (2021) [hereinafter Berkey, *Philosophical Core*].

11. Peter Singer, *Famine, Affluence, and Morality*, 1 PHIL. & PUB. AFFS. 229, 234 (1972).

excessively burdensome. If giving to the level of marginal utility is too big an ask, 10% is a friendly amount, one which would require little to no lifestyle change on the part of most of the affluent people EA targets. This is what MacAskill may be getting at in the interview I quoted above: 10% “is achievable yet at the same time really is a difference compared to what [people] otherwise would have been doing.”<sup>12</sup>

Some EAs are skeptical of this wariness about demands. In a paper on the value of “demanding the demanding,” Sachs points out that there’s limited empirical evidence that asking a lot of people “turns them off,” and that evidence does little to support EAs’ attitudes on this particular issue.<sup>13</sup> Instead, Sachs claims, the evidence points toward the “fudge factor” theory: the idea that most of us will live up to our commitments most of the time, only fudging a little every now and then.<sup>14</sup> If we demand the demanding of others, we might actually expect that they’ll do more than if we’d asked less of them.

If Sachs is right, then this takes the sting out of calling for EAs to give more than 10%. EA could be symmetrical: just as EAs are required to give to the *most effective* cause(s), they would be required to give *as much as they can*. They might fudge a little, but they’ll certainly give more than 10%.

The problem (as Sachs himself admits) is that there isn’t a whole lot of evidence to support the fudge-factor theory in this context.<sup>15</sup> The centerpiece of Sachs’s evidence is a paper examining how likely people are to cheat; the study authors found that when people are given opportunities to cheat on a test for minor financial bonuses, they consistently cheat, but only by a little bit.<sup>16</sup> We see, in other words, that people cheat even when the rewards are very small and even when the case is a one-off.

While Sachs thinks we can extrapolate from this case to EA, there’s an alternative, at least as plausible, interpretation: if people are willing to cheat for small (ten-cent!) gains in a one-off experimental condition, they may fudge even more, or outright refuse to comply, if presented with a much more demanding, and life-long, standard. Sachs does recognize that this is a possibility, but we should be pretty skeptical that we can draw broad conclusions about demanding the demanding from this study.<sup>17</sup> We don’t have good reason yet to be confident that the asymmetry is justified on these pragmatic grounds.

---

12. Stern, *supra* note 7. See also MacAskill et al., *Giving*, *supra* note 5.

13. Ben Sachs, *Demanding the Demanding*, in *EFFECTIVE ALTRUISM: PHILOSOPHICAL ISSUES* 138–43 (Hilary Greaves and Theron Pummer eds., 2019) [hereinafter Sachs].

14. *Id.* at 143–48.

15. *Id.* at 147.

16. Nina Mazar et al., *The Dishonesty of Honest People: A Theory of Self-Concept Maintenance*, 45 *J. MARK. RSCH.* 633 (2008).

17. Sachs, *supra* note 13, at 142–43.

### C. EA as a Project, not an Obligation

A second difference between EA and act utilitarianism, some EAs argue, is that EA functions as a discretionary project or an opportunity, something we can permissibly choose to take on (or not), rather than a moral obligation. MacAskill is one of the leading voices here: in his definition of EA, he casts the movement as a set of “two projects, rather than a set of normative claims”—the project of understanding how we can maximize the good we do, and the project of actually doing that good.<sup>18</sup> This definition overlaps with our discussion in the last section. If EA is just a project, it makes sense that it wouldn’t demand too much of us; moral obligations may lay claim to more of our resources. This project-based view of EA, in turn, would support the asymmetry: EAs may take on the project of effective giving, but they don’t have a moral obligation to give as much as they can.

There are a few issues here. First, MacAskill cites empirical data about EAs’ own views to argue that their movement is best understood as a project—but what does this data actually tell us? The first piece of evidence is a survey of EA “leaders” from 2015: 80% thought EA “should not include a sacrifice component.”<sup>19</sup> While obligations don’t necessarily involve sacrifices, discretionary projects are even less likely to, so this is a reasonable proxy for the project view. EA “leaders” and “members” disagree pretty starkly on this point, though. In 2017, a majority (over 56%) of EA members surveyed said they thought of EA more as a duty or obligation rather than an opportunity.<sup>20</sup> An earlier survey found that 34% thought EA was just an opportunity, while a total of 63% saw it as at least an obligation (some of that 63% believed it was both).<sup>21</sup> At the very least, then, there’s deep disagreement within EA about the nature of the movement—and if the survey of leaders was small-scale, and the surveys of members were much bigger, then there could in fact be a pretty wide consensus among EAs that their movement is tracking a moral obligation.<sup>22</sup> Moral obligations may not be all-consuming (the Kantian imperfect duty of beneficence is often thought to give us some latitude to decline to act<sup>23</sup>), but they don’t normally have percentage-based stopping points: it’s arbitrary to cap the obligation to give effectively at 10%.

18. MacAskill, *Definition*, *supra* note 5, at 14.

19. *Id.* at 15. There are a couple other aspects of the survey that are worth noting. 2015 is a long time ago in EA terms, and the composition and views of EA “leaders” may have changed pretty significantly since then. (One pair of relevant markers: 2015 is the year MacAskill’s global poverty-focused *Doing Good Better* came out; this year saw the publication of his longtermist *What We Owe the Future*.) MacAskill doesn’t provide a citation for this survey, and I haven’t been able to track it down.

20. McGeoch & Hurford, *supra* note 10; MacAskill, *Definition*, *supra* note 5, at 16.

21. *See* Cundy, *supra* note 10.

22. The 2015 survey had about 2300 respondents; the 2017 one had about 1800. *See id.* at 15–16. Since I wasn’t able to track down the 2015 survey of EA leaders, I can’t say for sure how strong the EA consensus is overall.

23. Karen Stohr, *Kantian Beneficence and the Problem of Obligatory Aid*, 8 J. MORAL PHIL. 45 (2011).

Maybe we shouldn't trust this data, because most EAs may not be familiar with philosophical terms of art or use them in the same way philosophers do. If this data is untrustworthy, this suggests we shouldn't rely too much on EAs' self-reports of their moral views; maybe they really do believe EA is a non-obligatory project, but they're misdescribing their own views. I'm not sure this gets us very far. The 2017 survey let participants respond that they either had no opinion or weren't familiar with different moral theories; non-philosophically inclined EAs had an out. Unless we have reason to think that a significant number of survey participants were mistaken about their familiarity with moral theory, we should assume that most (nearly 75%) of EAs have enough awareness to be able to pick a side. This near-unanimity is bolstered by the posts on EA Forum, which show a much higher degree of philosophical sophistication than we normally expect of non-philosophers.<sup>24</sup> People who post to EA Forum are a skewed sample even within EA; still, I think we'd need some significant evidence in order to be persuaded that the surveys have given a misleading impression of EAs' understanding of basic tenets of moral theory.

MacAskill's project-based view of EA isn't just a statistical generalization, though. He also argues that this is the right way for EA to be because it makes the movement more ecumenical. People with lots of different moral views can reasonably take up the non-moral project of learning how to give effectively. Giving What We Can found "anecdotal evidence" early on that people are more likely to pledge 10% when this is framed as an opportunity rather than as an obligation.<sup>25</sup> (If this anecdotal finding holds up, it is evidence against Sachs's argument in the previous section: perhaps "obligation" really does have a burdensome ring to it, and that's why GWWC's "opportunity" framing worked better.)

Yet, ecumenism has its limits. It might be nice in theory, but movements have to stand for something—and that means they cannot be completely ecumenical. If the point is to be effective, a movement which attracts and retains a small, highly committed core group may be better at that than a larger, diffuse movement whose minimal buy-in makes it easier to drift away over time.<sup>26</sup>

24. See richard\_ngo, *Scope-Sensitive Ethics: Capturing the Core Intuition Motivating Utilitarianism*, EFFECTIVE ALTRUISM F. (Jan. 15, 2021), <https://forum.effectivealtruism.org/s/aH5to3as8yiQA6wGo/p/NFGEgEaLbtyrZ9dX3> [<https://perma.cc/A84X-JMA3>]; Paul\_Christiano, *Integrity for Consequentialists*, EFFECTIVE ALTRUISM F. (Nov. 14, 2016), <https://forum.effectivealtruism.org/s/aH5to3as8yiQA6wGo/p/CfcvPBY9hdsenMHCr> [<https://perma.cc/A3N8-J4RA>]; MichaelaA, *Making Decisions Under Moral Uncertainty*, EFFECTIVE ALTRUISM F. (Jan. 1, 2020), <https://forum.effectivealtruism.org/posts/ex834aaANLhamLkvf/making-decisions-under-moral-uncertainty> [<https://perma.cc/SR6R-59GE>]; Aaron Gertler, *Utilitarianism With and Without Expected Utility*, EFFECTIVE ALTRUISM F. (Jul. 24, 2020), <https://forum.effectivealtruism.org/posts/BFDojdoTcBuqZyCp3/utilitarianism-with-and-without-expected-utility> [<https://perma.cc/2N5P-ZEGP>].

25. MacAskill, *Definition*, *supra* note 5, at 16.

26. See Owen Cotton-Barratt, *How Valuable Is Movement Growth?*, 1–21 (Center for Effective Altruism, Working Paper, 2015). I discuss this further at Amy Berg, *Effective Altruism: How Big Should the Tent Be?*, 32 PUB. AFFS. Q. 269, 275 (2018).

We need to know, then, whether a more ecumenical movement would be better at achieving EAs' goals and expressing their values. Berkey thinks it wouldn't be. If EA is a project, not an obligation, then it isn't wrong not to choose to give in a different way (or maybe even not at all), just as it wouldn't be wrong not to take up softball or cooking or bowling. And yet, Berkey points out, EAs criticize non-EA giving all the time.<sup>27</sup> When Singer criticizes David Geffen's multimillion-dollar donation to the Lincoln Center for the Performing Arts, Singer isn't merely politely pointing out that this giving isn't effective, nor is he saying that Geffen is failing to live up to some EA commitment he's made—Singer is saying that such giving is wrong, full stop.<sup>28</sup> Same thing with even seemingly more noble causes—see Singer's criticism of Batkid, or of guide dogs, or MacAskill's criticism of the ice bucket challenge.<sup>29</sup> If EA is just a project, these criticisms make little sense; you cannot criticize people who simply fail to share your own discretionary projects.

EAs could maintain the asymmetry by giving up their practice of criticizing non-EA charitable giving. This, though, would chip away at some of what binds them together as a social movement. Movements are generally predicated on some common cause, advocating for some way the world ought to change, and EAs have historically had a clear view of what change should happen: people are giving their money to the wrong things when they could be doing so much more good. A movement that says "We're donating here, but it's fine if you want to give there, or maybe not give at all" doesn't fit with the attitude the EAs I know have toward their own movement and the pitch they make to prospective members: "We know something you don't about how you should be living; do you want to find out what that is?" Appeals to join EA that start from the premise that EA is merely an optional project feel much limper.

At the same time, social movements of all stripes have lots of internal disagreements and yet seem able to hang together reasonably well as movements. Could there be a project and an obligation wing of EA, where the first is fine with

27. Berkey, *Philosophical Core*, *supra* note 10, at 103–04.

28. Singer has criticized this donation in a number of places, including in a *Wall Street Journal* interview. Alexandra Wolfe, *Peter Singer on the Ethics of Philanthropy*, WALL ST. J. (Apr. 3, 2015), <https://www.wsj.com/articles/peter-singer-on-the-ethics-of-philanthropy-1428083293> [<https://perma.cc/TRA7-R5HY>].

29. Singer talks about Batkid in a *Washington Post* article from 2013. Peter Singer, *Heartwarming Causes Are Nice, but Let's Give to Charity with Our Heads*, WASHINGTON POST (Dec. 19, 2013), [https://www.washingtonpost.com/opinions/heartwarming-causes-are-nice-but-lets-give-to-charity-with-our-heads/2013/12/19/43469ae0-6731-11e3-a0b9-249bbb34602c\\_story.html](https://www.washingtonpost.com/opinions/heartwarming-causes-are-nice-but-lets-give-to-charity-with-our-heads/2013/12/19/43469ae0-6731-11e3-a0b9-249bbb34602c_story.html) [<https://perma.cc/D26M-R2F9>]. He makes the guide-dog critique often. See, e.g., Rachel Elizabeth Maley, *Is There Such a Thing as Bad Charity?*, THE LIFE YOU CAN SAVE (May 24, 2015), <https://www.thelifeyoucansave.org/blog/is-there-such-a-thing-as-bad-charity/> [<https://perma.cc/BZA7-ZWRS>]. MacAskill critiques the ice bucket challenge on at least two occasions. See William MacAskill, *This Week, Let's Dump a Few Ice Buckets to Wipe Out Malaria Too*, QUARTZ (Aug. 18, 2014), <https://qz.com/250845/this-week-lets-dump-a-few-ice-buckets-to-wipe-out-malaria-too> [<https://perma.cc/RAB9-PZ4H>]; William MacAskill, *The Cold, Hard Truth About the Ice Bucket Challenge*, QUARTZ (Aug. 14, 2014), <https://qz.com/249649/the-cold-hard-truth-about-the-ice-bucket-challenge> [<https://perma.cc/Z3UJ-KYSL>].

donations like Geffen's and the second is critical of them? If so, the asymmetry would only really bite for the obligation-focused wing; the other wing could say that there's no problem with choosing not to take on the project of giving 10%. At the end of this paper, we'll discuss some of EA's other core commitments—it seems unlikely to me that project-based EA is congruent enough with these commitments to count as true EA, but we can keep this open as a possibility for now.

In the meantime, there's a third problem with a project-based view of EA, something that should lead us to suspect this wing of the movement cannot get off the ground. It isn't as though EAs think their project is *only* about researching effectiveness (although more on this in a bit); after all, they do seem to think that EA commits them to giving at something like the 10% level many EA organizations emphasize. But if EA is just a project, then where does 10% come from? We're normally morally and prudentially permitted to commit to our discretionary projects at whatever level we think is best for us. I could take up pottery as a career, as a serious hobby, or as something I do only every few months. Any of these choices seem to be fine—it is just a question of what's best for me. Why can't EA be a project a person decides to spend only 5% (or 1% or 75%) of her earnings on?

Some projects do come with moral strings attached. If you're a parent, you have moral obligations to your children; if you're on a kickball team, you owe it to your teammates to show up for games. (Compare this to EA as conditional obligation in the section after next.) Perhaps something like this applies even to a project-based view of EA; if you've pledged to give 10%, you have some obligation to keep your word and not let your fellow EAs down. But this doesn't explain why 10% is the pledge amount in the first place; this is still arbitrary.

It could also be that some people take on the project of EA for partly moral reasons, just as someone might believe they're morally obligated to become a parent or share their novel with the world. What separates the project view of EA from the obligation view of EA is not whether someone could believe they have moral reasons for engaging in a particular project but rather whether EAs should say EA is an obligation we incur in general. If EA is a project some people could adopt for moral reasons but that no one is required to take on, 10% still lacks a good justification.

For these reasons, EAs should reject the project view of EA. That, in turn, casts doubt on the asymmetry. Moral obligations don't allow us to commit ourselves at arbitrary levels we set at our own discretion. Not all moral obligations are maximizing, but they also aren't only sensitive to the kinds of pragmatic or prudential concerns we use to decide how much time to devote to our different life projects.

One complication is whether EA's new longtermist focus affects how we think of the movement. It's significant that MacAskill's recent work on longtermism avoids discussion of obligation almost completely. Despite its title, *What We Owe the Future* is framed around showing us that we have the *opportunity* to shape the far future, and that it would be good if we took advantage of this opportunity, but MacAskill stays away almost completely from claiming that we ought

to do so.<sup>30</sup> This is particularly notable because MacAskill was one of the founders of Giving What We Can and one of the major advocates of the 10% standard; his 2015 *Doing Good Better*, which was focused primarily on global poverty, made a number of claims about what we should do.<sup>31</sup> (MacAskill does make normative claims about morality in some recent work, and some of these claims are potentially relevant to EA, but he steers away from applying those claims directly to EA as a movement.)<sup>32</sup>

It's not clear that longtermism is necessarily a nonnormative view; Ord's *The Precipice*, which anticipates many of the moves in *What We Owe the Future*, frames longtermism as an "ethic" and connects it to "a great many schools of thought about morality."<sup>33</sup> Even so, Ord frames his suggestions primarily in terms of what we *can* do, not what we *should*.<sup>34</sup> Overall, longtermist EA focuses almost exclusively on effectiveness, and hardly at all on altruism—the asymmetry is only going to widen as longtermism's influence grows. This could fit the project-based view of EA, yet we've seen that this view may be at odds with some of the commitments binding EA together as a movement.

#### D. Research Priorities

And yet it may be what MacAskill, at least, has in mind. If we see EA as non-normative, he writes, then this

focuses attention on the most distinct aspect of effective altruism: the open question of how we can use resources to improve the world as much as possible. This question is much more neglected and arguably more important than the question of how much and in what form altruism is required of one.<sup>35</sup>

---

30. In Appendix 4, the only place in the book where MacAskill really broaches this question, he moves away from it almost immediately:

How much *should* we in the present be willing to sacrifice for future generations? I don't know the answer to this. All I've claimed in this book is that concern for the longterm future is at least one key priority of our time . . . . But it does seem to me that we should be doing much more to benefit future generations than we currently are.

WILLIAM MACASKILL, *WHAT WE OWE THE FUTURE* 261 (2022) [hereinafter MACASKILL, *WHAT WE OWE*].

31. MACASKILL, *DOING GOOD BETTER*, *supra* note 8. Among them: We should ask whether we are putting our money to its *best* use, *id.* at 88–89, we should not donate on the basis of personal connections, because this is unfair, and indeed donating to less effective charities is "absurd," *id.* at 40, we should resist the inclination to give based on our emotions, *id.* at 60–61. While it might be that MacAskill was thinking of these as hypothetical imperatives for people already committed to EA, he wrote the book to convince people to accept EA—it seems clear he's trying to persuade his audience that they have independent moral reason to become EAs.

32. He and Mogenson argue that both nonconsequentialists and consequentialists are morally committed to being longtermists. Andreas Mogensen & William MacAskill, *The Paralysis Argument*, 21 *PHILOSOPHERS' IMPRINT* 1, 15 (2021).

33. TOBY ORD, *THE PRECIPICE: EXISTENTIAL RISK AND THE FUTURE OF HUMANITY* 84, 101 (2020).

34. *Id.* at 362.

35. MacAskill, *Definition*, *supra* note 5, at 16.

In other words, EAs should emphasize the effectiveness side of the asymmetry, because the altruism side—the question of how much to do—has historically gotten more attention.

EA can't completely ignore the altruism side, and not all EAs are on board with MacAskill's approach. Pummer makes almost exactly the opposite point: "Finding out how to help others the most seems to be part of the project of effective altruism only if it is part of helping others the most."<sup>36</sup> EA may be distinctively focused on research, but if that research isn't put to use to do good, it is useless. A highly research-focused movement could be even more niche, less ecumenical, and less effective. It would be a bad thing if EA as a whole took on the character of some of its more abstruse online corners, busily debating how many QALYs can dance on the head of a pin. (Again, there are parallels to utilitarianism here—utilitarians have to decide when continuing to determine what maximizes utility detracts from their ability to actually maximize utility.)

Indeed, although we know MacAskill positions the project view of EA as ecumenical, focusing on the research side of this project could backfire. Nearly anyone who's in a position to have heard of EA is in a position to donate some of their earnings, and most of those people are comfortable enough that they could give 10% without its being a major sacrifice for them. Many fewer of these people are in a position to run a randomized controlled trial.

This is particularly notable since research-focused EA is on the rise. We already saw that longtermists shy away from discussion of moral obligations. Instead, longtermist EA focuses on making sure that the people who are capable of doing research into longterm threats—AI, biological weapons, climate change, nuclear war—can do so productively. For the rest of us, altruism drops out almost entirely. MacAskill's suggestions for how to be a longtermist boil down to donating to effective nonprofits, being politically active, discussing longtermist ideas with friends and family, and having children.<sup>37</sup>

Even donation seems to be less of an emphasis than it was for global poverty-focused EA. The only longtermist organization I can find focused on helping people decide where to donate is Longview Philanthropy, which "designs and executes bespoke giving strategies for major donors," those who are in a position to give away at least a million dollars a year.<sup>38</sup> The transparency and original pitch of early EA—that you, yes you!, can save lives by donating relatively small

---

36. Similarly, he writes, ". . . it seems you are engaged in effective altruism if *relative to your evidence* you help others the most, even if *in fact* you fail to help others at all." PUMMER, RESCUE *supra* note 5, at 208–09 n.7.

37. MACASKILL, WHAT WE OWE, *supra* note 30, at 232–34.

38. MacAskill is an advisor to Longview. Since I originally wrote this article, Longview Philanthropy has partnered with Giving What We Can to create a Longtermism Fund, which accepts small contributions. See Michael Townsend, *Longview Philanthropy: Longtermism Fund*, <https://www.givingwhatwecan.org/charities/longtermism-fund> [<https://perma.cc/P6BN-S4NB>]. This is progress, although the Longview Philanthropy homepage still emphasizes donors "giving at least \$1 million . . . per year," and the small-donor option is difficult to find unless you know what you're looking for. *Id.*

amounts—are relatively absent from longtermist EA, supplanted by a greater focus on people who have serious money to throw around.

It may not be unreasonable to expect longtermist EA to focus more on research, since we know relatively more about the poor who are alive today than we do about the fantastical possibilities—space colonization, artificial general intelligence, civilizational collapse—the longtermists present us with. Research-focused, project-based, longtermist EA can make sense of the asymmetry between the two halves of the movement, assuming that research should be the overwhelming priority of the movement. This comes at the significant cost, though, of excluding many current EAs from playing a meaningful role in their own movement and of moving away almost entirely from the altruism side.

### *E. EA as a Conditional Obligation*

Some EAs disagree with MacAskill that EA is just an opportunity. There are obligations to give effectively, they argue, but only *conditional* obligations: *if* you've taken on the project of helping others, *then* you have an obligation to give effectively. Pummer is a proponent of this approach to EA, arguing from cases such as:

*Arm Donor:* There is one innocent stranger stuck on track A, and one hundred strangers stuck on track B. On each track is a trolley headed toward them, which will kill them unless the trolley is stopped. You can use your arm to stop one, but not both, of the trolleys, and will lose the arm if you do.<sup>39</sup>

Pummer claims that you have only a conditional obligation here: you're not obligated to sacrifice your arm, but if you're willing to make that sacrifice, then you're obligated to save the hundred; it would be wrong for you to save the one instead.<sup>40</sup> Likewise, it isn't wrong to refuse to donate your earnings when doing so is costly for you, but once you've decided to use some of your earnings to help others, you must do so in the most effective way.

Seeing EA as a conditional obligation solves a few of the problems with the project-based view of EA we discussed above. It fits with EAs' own conception of their movement as more than a project.<sup>41</sup> It appears to be able to explain other elements of EA practice, such as criticizing some charitable giving. Once Geffen decides to give away parts of his fortune to help others, he's taken on the conditional obligation to do so effectively, and he can be criticized for failing to live up to that requirement. Yet conditional EA has the potential to be more ecumenical

39. This paraphrase is from Berkey, *Effectiveness*, *supra* note 5, at 373. The original case appears in Pummer. See Theron Pummer, *Whether and Where to Give*, 44 PHIL. & PUB. AFFS. 77, 83 (2016) [hereinafter Pummer, *Whether and Where*]. Horton presents similar cases. See Joe Horton, *The All or Nothing Problem*, 114 J. PHIL. 94, 94 (2017). Although he doesn't tie them directly to EA in the way Pummer does, the implications for would-be EAs are pretty clear. See PUMMER, *supra* note 36, at 41.

40. See Pummer, *Whether and Where*, *supra* note 39, at 83–84.

41. MacAskill notes the conditional view of EA as a possibility. See MacAskill, *supra* note 19, at 17.

than a stricter, unconditional EA would be, since we aren't required to give when doing so is costly to us.

If it's a workable version of EA, then, conditional EA is poised to justify the asymmetry. When we choose to give, we're required to give effectively—so a clear, objective, maximizing standard for effectiveness is justified. We may have some unconditional obligations to give, when doing so isn't costly to us—Pummer (2023, 201) uses 10% as the ballpark figure, although he acknowledges this could be lower or higher. Beyond that obligatory amount, whether we give when it's costly is up to us.

Conditional EA faces a dilemma, though: it's too consequentialist to be truly ecumenical, yet it's not consequentialist enough to satisfy consequentialist EAs. Sinclair objects that conditional EA displays a “halfhearted” nonconsequentialism at best. Conditional EA gives us options, but, Sinclair writes,

moral options are effectively conceived as a domain of licensed non-optimific action against a background presumption that action ought to be optimific. This is consequentialist in broad outlook, even though it allows for the presumption's defeat in cases of excessive cost to the agent.<sup>42</sup>

There are certainly EAs who accept this halfhearted nonconsequentialism; Pummer is one. But many other (what Sinclair calls “thoroughgoing”) nonconsequentialists, those who reject the idea of a “default presumption of an optimizing or burden-minimizing standard,” do not. Thoroughgoing nonconsequentialists don't have to accept the requirement, bedrock to EA, that we maximize our effectiveness when we give.

Conditional EA takes fire from the other side too. Since Pummer allows that people may give to less effective charities whenever they care more about the work the less effective charity is doing,<sup>43</sup> this limits EA “quite a bit more than is suggested by much EA advocacy and argument.”<sup>44</sup> If, for whatever reason, you care more about the work a guide-dog nonprofit does than about the work a trachoma-prevention NGO does, then you seem to be permitted to give to the less effective organization. Pummer can avoid problems like this by strengthening his account of conditional obligations: “It is wrong to perform an act that is *much worse* than another, if it is *slightly* costlier to you to perform the better act, and if all other things are equal.”<sup>45</sup> But, Berkey says, now unconditional obligations have reentered the picture: whenever it is much better to give to an effective charity at a slight cost to ourselves, we have the *unconditional* obligation to do so.<sup>46</sup> If EAs want to retain the core of their movement, they “cannot avoid taking at least

42. Thomas Sinclair, *Are We Conditionally Obligated to Be Effective Altruists?*, 46 PHIL. & PUB. AFFS. 44 (2018) [hereinafter Sinclair].

43. See Pummer, *Whether and Where*, *supra* note 39, at 82.

44. Berkey, *Effectiveness*, *supra* note 5, at 377.

45. See Pummer, *Whether and Where*, *supra* note 39, at 93.

46. Berkey, *Effectiveness*, *supra* note 5, at 378.

a fairly strong position on morality’s demandingness,” a position which excludes thoroughgoing nonconsequentialists.<sup>47</sup>

Conditional EA may support the asymmetry, but otherwise it’s neither fish nor fowl. As a view that we must maximize the impartial good unless specifically permitted not to, it isn’t ecumenical enough for most nonconsequentialists. As a view that permits significant departures from what’s impartially best, it is at odds with the views of most consequentialists (and most EA theory and practice). There’s a thin slice of EAs (the “halfhearted” nonconsequentialists Sinclair describes) who may be comfortable with this, but EA as a conditional obligation seems to make the view less, rather than more, ecumenical.

This is not to say that EAs don’t have possible avenues of response. Some conditional obligations (such as the obligation of gratitude) are part of the moral landscape; maybe there are ways to reframe EA so it resembles those more widely accepted conditional obligations.<sup>48</sup> The question is whether this reframing can also support the ecumenism that was one major reason for turning to conditional obligations in the first place.

#### F. Alienation

EA might set a 10% standard because an alternative, higher standard—20%, 50%, or more—would mean people could no longer reserve time and money for their own lives or act in accordance with their own integrity, and this would be alienating.<sup>49</sup> This supports the asymmetry—people aren’t alienated by giving effectively, but they are alienated by overdemanding altruism.

It does seem potentially alienating to give a lot of your money away, but this can’t explain the asymmetry. Krishna writes, “At once pious and rational, comforting and selfless, effective altruism promises a life free from all the hokeyness involved in the business of finding ourselves and our deepest impulses.”<sup>50</sup> If you don’t have any causes you care about, he seems to be saying, EA can take care of that, and you don’t even have to do much reflecting on anything. If, on the other hand, you already have things that are important to you, EA tears you away from what really matters to you. Your child had a traumatic experience with cancer, say, and you don’t feel alienated when you give money to researching neglected childhood cancers—but you do feel alienated when you give to schistosomiasis control.

47. *Id.* at 380. This may not be unwelcome news to all EAs; it’s not clear from the survey data that EAs who accept obligatory EA think that this obligation is merely conditional. The majority of EAs who are act utilitarians may believe that the suffering of others gives us unconditional obligations, even when this is costly to us.

48. Sinclair, *supra* note 42, at 38–39.

49. Once again, this is an objection to consequentialist theories that also seems to work against EA. For more on consequentialism and alienation, see generally Peter Railton, *Alienation, Consequentialism, and the Demands of Morality*, 13 PHIL. & PUB. AFFS. 134 (1984).

50. Nakul Krishna, *Add Your Own Egg*, THE POINT (Jan. 13, 2016), <https://thepointmag.com/examined-life/add-your-own-egg/> [<https://perma.cc/2D2V-R9K6>].

Why, though, should we think that giving 10% of your earnings away to schistosomiasis control would be less alienating than giving 10%, or even 20%, of your earnings to cancer research? People frequently spend huge amounts of time and money on aiding others even when this is costly. Good parents spend significant chunks, often much more than 10%, of their earnings on their kids, with no expectation of return and not primarily because it benefits them personally. Yet this doesn't seem to be alienating (in general—of course there are exceptions, even among good parents), because this is a project parents themselves choose to engage in. Alienation, then, doesn't support the asymmetry: it is alienating to give too much away, but it is also alienating, possibly more so, to give even relatively small amounts to causes you're disconnected from.

EAs argue that it is possible to decrease the alienation we feel when giving. In general, MacAskill, Mogenson, and Ord write, we tend to overestimate the duration and intensity of the unpleasant feelings we have when we experience bad outcomes, so we might overestimate how alienated we'll be if we give to causes with which we feel no connection.<sup>51</sup> MacAskill reflects on his own experience: "I feel that my life is more meaningful, authentic, and autonomous . . . Effective altruism has added to my life, not subtracted from it."<sup>52</sup> But that's just one guy. Bowen thinks that this generalizes to the rest of us; he writes, "If we had access to better information and took time to reflect on how we choose between charities, I expect most people would realize that what they actually care about is improving the lives of as many people as they possibly can by as much as they possibly can."<sup>53</sup> We should ask ourselves questions like, "If a loved one of yours were to die from cancer, would this make children who die from malaria less important than children who die from cancer?" and if we do, he concludes, "I imagine most people would realize how arbitrary and unimportant factors like physical proximity are to them."<sup>54</sup>

Bowen cites some empirical work that appears to lend some support to this claim—research shows people do care about cost-effectiveness<sup>55</sup>)—but he's extrapolating from that to the idea that people could come to think like EAs if they just thought through their own commitments rationally enough, and he lacks evidence to make that leap. It's easy to imagine someone responding: "Sure, I agree that children with cancer aren't more important than children with malaria. But children with cancer are more important *to me*, because I know the kind of suffering my child went through when she was battling cancer, and since I can't solve every problem, I want to focus on the problem that I know and care about." This may not be the morally *right* response, but it seems like how many of us actually do think about these things. It might be great if we were totally rational

---

51. MacAskill et al., *Giving*, *supra* note 5, at 184.

52. MACASKILL, *WHAT WE OWE*, *supra* note 30.

53. Dillon Bowen, *The Economics of Morality*, 4 J. PRAC. ETHICS 80, 94 (2016).

54. *Id.* at 95.

55. *Id.* at 94.

utility calculators, but we're not, and just asking probing questions about our own values is not enough to dissolve the alienation worry.

This isn't to say that Bowen's wrong about cost-effectiveness. It does seem right that people prefer to give to more rather than less effective groups working on *the things they care about* (for example, preferring to give to a more rather than less effective pediatric-cancer charity). But it's a big leap to conclude that people would prefer to give to more rather than less effective groups *regardless of cause*. Without more evidence, we shouldn't rush to conclude that people can get rid of the alienation they feel just by thinking through their own commitments. Many people may find EA-style giving alienating even when they give relatively low amounts.

This is another way longtermism is poised to disrupt EA. As longtermism becomes a more significant part of EA, I'll be curious to see whether the movement itself comes to alienate some of its members. Will people who came to a movement to do something about global poverty be turned off by the view of some longtermist EAs that global poverty is a "rounding error" when compared to the long term?<sup>56</sup> Or will they update their own priorities as EA's focus changes? We may find out.

### III. RESOLVING THE ASYMMETRY

We began with an asymmetry—that effective altruists have stringent criteria for effectiveness, but they're relatively lax about how much altruism is required. We then looked at several possible explanations for this asymmetry. EAs might say that we can avoid arbitrariness in where to give—but that doesn't fully hold up in the face of disagreement about effectiveness and a relatively clear act-utilitarian standard of beneficence. The asymmetry may be justified for pragmatic branding reasons—people don't want to join a movement that's too demanding—but empirical evidence for this is minimal, and even some EAs think that EA should go ahead and demand the demanding. EA might be a project, rather than an obligation—but whether EAs actually believe this is questionable, and if EA is a discretionary project, then that calls for an even laxer standard of giving. Research-focused EA may be more ecumenical, but it seems like it resolves the asymmetry by nearly completely bypassing the altruism side of EA. EA might instead be a conditional obligation, but this winds up being even less ecumenical than other options. EA that requires people to give too much may be alienating, but it isn't clear why it would be more alienating than EA that requires people to give to a circumscribed set of causes.

---

56. The "rounding error" view was described in a 2015 Vox article by Dylan Matthews (himself an EA). See Dylan Matthews, *I Spent a Weekend at Google Talking With Nerds About Charity. I Came Away . . . Worried*, VOX (Aug. 10, 2015, 10:00 AM), <https://www.vox.com/2015/8/10/9124145/effective-altruism-global-ai> [<https://perma.cc/MW85-ZGZP>]. This idea isn't new, but it's more influential than it used to be.

This doesn't prove conclusively that the asymmetry is unjustified, but it should make us skeptical. Unless and until there's some other explanation justifying the asymmetry, EAs should bring the two halves together: either by raising the standard of altruism or by lowering the standard of effectiveness.

Start with the first option. In this case, EAs apply the same stringent standards to how much to give as they do to where to give. Just as giving to the Malaria Consortium allows us to do more good than giving to Oxfam or Heifer International, so too giving to the level of marginal utility allows us to do much more good than giving just 10% would. Of course, EAs acknowledge room for disagreement on issues like the relative importance of animal suffering, so maybe a symmetrical EA ought to have some space for disagreement about how much to give too. But just as EAs agree that something like QALYs is the right measure of effectiveness, and Geffen's gift to Lincoln Center can in no way count as effective, so they can only brook so much disagreement about how much you ought to give. This wouldn't bring EA all the way back around to act utilitarianism, but it would bring it closer to the demands—and controversies—of that theory than some of its proponents might like to admit.

The other option is to lower the standard for how effective giving must be to count as EA. Just as EAs have a relatively minimal standard for how much you must give to count as EA, and they accept a wide range of giving beyond that threshold, so too they could have a relatively minimal standard for effectiveness and a wide range of "effective" options beyond that. It's hard to say what that relatively minimal standard might look like—does something to reduce suffering? Doesn't actively waste money?—but, if it's going to parallel effectiveness, it has to be pretty minimal. Just as 10% is supposed to be so easy to give away that we might not even notice it, a comparable effectiveness requirement would have to be pretty lenient.<sup>57</sup>

Yet reducing the standard of effectiveness in this way violates core EA commitments. There are different, although overlapping, accounts of what exactly those core commitments are. Pummer defines EA as "the project of using time, money, and other resources to help others the most."<sup>58</sup> The Center for Effective Altruism says it's "using evidence and reason to figure out how to benefit others as much as possible and taking action on that basis."<sup>59</sup> Based on definitions like these, I've argued elsewhere that EA's most central commitment is that it provides a single objective standard by which we can analyze our giving. This standard allows for critical scrutiny of giving, allows EAs to avoid various psychological phenomena they see as problematic (such as confirmation bias or emotional attachment to a pet cause), and commits EAs to doing good in perpetuity.<sup>60</sup> Berkey gives a slightly longer list of core commitments: 1) strong interest-

---

57. MacAskill et al., *Giving*, *supra* note 5, at 181–89.

58. See PUMMER, *RESCUE* *supra* note 5, at 181.

59. MacAskill, *Definition*, *supra* note 5, at 13.

60. Berg, *supra* note 26, at 269–87.

based reasons (the claim that others' interests give us strong reasons to meet their needs), 2) cosmopolitan impartiality (the claim that we should direct our charitable resources wherever they can do the most impartial good), and 3) evidence-based decision-making (that we should gather and use evidence about what will do the most good).<sup>61</sup>

EA with lax standards for both effectiveness and altruism pretty clearly violates most of those core claims. It violates the maximizing elements of both Pummer's and the CEA's definitions, since EAs are no longer centrally concerned with maximizing the good they do. It's also at odds with the idea of a single objective standard for giving that's the same for everyone. Losing that single standard means losing the grounds for criticizing giving due to common psychological phenomena rather than to rational responses to empirical evidence. If people are allowed to give to even minimally effective causes, then emotional attachments, biases, and so on could dictate where they give. And it violates cosmopolitan impartiality (since it no longer directs its resources to where they do the most impartial good) and possibly also evidence-based decision-making (since decisions about where to give are now not made solely on the basis of the evidence about what's the most effective).<sup>62</sup>

But perhaps there are ways to both act effectively and give to the things you care about. Pummer argues that it is sometimes permissible (for example, on the basis of a personal connection) to give twice as much to a charity that's half as effective—call this *big altruism*.<sup>63</sup> Pummer uses 10% as the baseline requirement if you're giving as effectively as possible, although he claims that “for many of us, the relevant baseline lifetime requirement will not be 10 percent, but something considerably lower or higher.”<sup>64</sup>

At base, big altruism accepts the asymmetry: you have to be as *effective* as you would be if you were a standard EA, and it is the *altruism* (that is, the amount of your giving) that flexes up or down depending on the precise charities you give to. In practice, it stands to be a reasonable compromise between raising the standard for how much to give and lowering the standard for how effectively to give. Big altruism has other advantages too. It is more precise than the blanket 10% standard, recognizing that some people are obligated to start from a baseline that's much higher than 10% of their earnings and that others would have to sacrifice a lot to give so much. It also avoids some of the alienation worries we discussed above, since it lets you be an EA in ways that fit with your own projects, commitments, and values.

61. Berkey, *Philosophical Core*, *supra* note 10, at 95–97.

62. Even project-based EA is in tension with these core commitments. EAs don't generally say that bias or giving based on sloppy evidence is bad only for their specific, voluntarily assumed projects; they seem to think that these are bad in general, regardless of whether you're a self-identified EA or not.

63. See PUMMER, *RESCUE* *supra* note 5, at 201–02. This was Pummer's own name for it in an early draft of his, and I think it's a helpful label. Dan Shahar also suggested something like this to me in personal communication years ago.

64. *Id.* at 201.

At the same time, there are limits to how much “big altruism” can avoid alienation. It’s hard to imagine an arts charity that could be ten times as good at adding QALYs to the world as the Malaria Consortium is—so if your obligation were supposed to have the effect on the world you would if you gave 10% of your earnings to EA causes, and you wanted an EA justification for giving to the arts, you’d have to sacrifice well over 100% (10x10%) of your earnings. We can’t know how many charities are permitted on Pummer’s style of EA, but it has to be a relatively small list.

Pummer’s view has at least one more theoretical cost. In original EA, you commit to do *the most good you can do* with the 10% you give. On Pummer’s view, you’re permitted to no longer have that commitment, at least not in a direct way. You’re taking steps to ensure that your overall effect on the world is the same, and maybe this indirect commitment to the spirit of EA is enough to count you in. Still, something fundamental to EA (something along the lines of Berkey’s “cosmopolitan impartiality”) seems to have been lost. Wouldn’t an orthodox EA look at you and say: you say you’re committed to doing good, so why not recognize that you can put those same resources to work and do even more good? Big altruism may be significantly costly to us, particularly if the causes we wish to give to aren’t especially effective, and at the same time it’s at odds with at least one core value of EA.

The 10% tithe found throughout much EA work, and the contrast it has with EAs’ emphasis on effectiveness, looks pretty shaky. And, given EAs’ own commitments, the best way to get rid of this unjustified asymmetry is to bring the laxer altruism standard in line with the more uncompromising effectiveness standard, rather than the reverse—and middle-ground attempts like big altruism aren’t yet workable compromises. Even so, that doesn’t solve all the moral questions around EA. The movement is more demanding than EAs like MacAskill claim, but it’s still not fully clear what the shape of those moral obligations is. Perhaps EAs have to go all the way to embracing act utilitarianism, or maybe there’s some more pluralistic standard they could adopt. Whatever this standard turns out to be, EA will undoubtedly make more demands than some of its proponents seem to think.