

**Hilary Greaves and Theron Pummer**, *Effective Altruism: Philosophical Issues* (Oxford, Oxford University Press, 2019), pp. x + 247.

Effective altruism (EA) is a movement with its roots firmly planted in academic philosophy. While it's escaped the bounds of academia in the last few years—the annual EA Global conference attracts such luminaries as Elon Musk—it's a young movement, still working out what it should stand for. Hilary Greaves and Theron Pummer's new edited volume, *Effective Altruism: Philosophical Issues*, is a welcome contribution to the still-developing theory behind EA. It gives a broad overview of EA at present and, in some of the themes which emerge over the course of the volume, of what might capture the attention of the movement in the future. The 15 essays in this volume address both core and peripheral theoretical issues which are of concern to EAs, philosophers working on EA, and those interested in issues of beneficence more broadly. Anyone who belongs to any of these categories will find the volume highly worth reading.

That's a real strength of this volume—there are lots of audiences who will potentially benefit from reading some of the contributions. On the other hand, some contributions will be more relevant to some audiences than others, which is probably par for the course in this kind of edited volume. If you're wondering whether you'd benefit from reading this volume, consider whether you fall into one of these camps:

You might want to read this volume if you're only vaguely familiar with EA, and you want to know what it's all about. You'd probably learn a lot from the contributions by MacAskill, Ord, and perhaps Beckstead. But I think you would find some of the later chapters rougher sledding. Some of them are better seen as contributions to debates about the duty of beneficence than debates about EA; a few of the later chapters scarcely mention EA. (This might be less of a concern if EA is a movement which makes normative claims about beneficence—more on this later.)

You might read this volume because you're familiar with EA and want a comprehensive overview of current philosophical debates, or maybe you want to use this volume to teach a class on EA. You'd benefit from reading the chapters by Askill, Snowden, Gabriel and McElwee, among others. (If you're in this camp, I would especially recommend Budolfson and Spears's contribution, which I thought was excellent.) At the same time, this volume isn't a completely comprehensive overview of EA. My sense is that EAs have become increasingly concerned with issues of animal welfare, and so I was surprised to see that no chapter was entirely or even primarily dedicated to animal ethics (although existential risk, another perennial EA concern, gets an outing in Beckstead's chapter).

I also think a comprehensive overview of EA would not be complete without an in-depth look at some critiques of the movement, and here I think the volume missed an opportunity. Gabriel and McElwee contend that EAs have not so far done enough to support systemic change, but this argument (if correct) serves as an internal critique of how EAs pursue their priorities rather than an external

critique of whether those priorities are good ones in the first place. Saunders-Hastings does offer an external critique—EAs, she argues, engage in a *pro tanto* problematic form of paternalism whenever they give to organizations that constrain recipients' choice (for example, by buying malaria bednets rather than making direct cash transfers to poor people). At the same time, this looks an awful lot like an indictment of most charitable giving, not just a critique of EA, although Saunders-Hastings does give some reasons to think that this may be an especially pressing problem for EAs (p. 123). A volume that could truly serve as a comprehensive overview of the pros and cons of this movement would have to contend with more, and more EA-specific, objections.

You might also read this volume looking for connections between EA and other philosophical debates. Here, you might look to the contributions by Sebo and Paul, Timmerman, and Collins, among others. At the same time, some of these issues seem (to me at least) to be at the periphery of philosophical debates around EA.

Or are they? That really depends on what EA's core commitments are, and that's a question EAs haven't fully settled yet. In the first chapter, MacAskill claims that EA is 'non-normative': it 'consists of two projects [to figure out how to do as much good as possible and to use this information to do good], rather than a set of normative claims' (p. 15). If this is true, then there is no obligation to be an EA, and if one does join the movement, then one is not obligated to give away as much as one can—but a commitment to EA does mean a commitment to using (some amount of) one's resources as efficiently as possible. MacAskill gives three principal grounds for treating EA as a non-normative project (pp. 15-16). In a 2015 survey of *leaders* of EA, 80 percent indicated they saw it as a non-normative project; 12.5 percent said the opposite. In a 2017 survey of *members* of EA, 56.5 percent indicated that they saw EA as a 'duty' or 'obligation' and 37.7 percent as an 'opportunity' (in an earlier survey, a plurality said it was both obligation and opportunity). Finally, if EA is a non-normative project, it can be a more ecumenical movement, consistent with many different moral views.

But take a look at these mixed results. While the vast majority of *leaders* of EA think EA should be a non-normative project, EA *members* seem much more inclined to say that EA is morally required. And while ecumenism is good, movements have to stand for something. As Gabriel and McElwee note, while 'effective altruism purports to be a movement that combines altruistic intentions and efforts with careful thought and attention to evidence,' 'in practice, effective altruists tend to adhere to a thicker set of norms and assumptions,' in particular welfarism and consequentialism (p. 100). Should EA abandon its aspirations to ecumenism?

That tension is brought out in a number of the contributions to this volume. Sachs's chapter takes this on directly: he contends that non-normative EA is motivated by a fear that it would be counterproductive to push people to make very large sacrifices (p. 137). The language of 'opportunity' doesn't demand sacrifice, but the language of 'obligation' might. And so if Sachs is right that the psychological

evidence doesn't show that 'demanding the demanding' would be counterproductive (he cites a very small number of studies in support of this claim, but he accepts that the psychology here is far from settled), and if he's right about EAs' implicit justification for not asking people to sacrifice, then perhaps EA will come to shed its non-normative guise.

If Sachs is right, and EA should demand the demanding, then this explains the inclusion of many of the chapters in the second half of the book. Barry and Lawford-Smith argue that determining whether someone has satisfied her duty to assist requires considering her 'inputs' (time, energy, and so on), 'characteristics' (intention to assist), and 'success' (at least some degree of efficiency in the means taken to assist) (p. 163). This is a debate about morality—what are the limits of my duty of beneficence? Likewise, Mogensen considers the Callousness Objection to Singer's famous pond case; this objection purports to show that there are some situations in which it would be callous to do the utility-maximizing thing rather than meet a salient nearby need (e.g., by saving a child drowning in a pond). This too is a debate about morality—what am I morally required to do when opportunities to act beneficently conflict? For these contributions to tell us something about EA, EA must make normative claims about the duty of beneficence—if EAs are merely taking advantage of an opportunity or committing to a certain personal project, then the limits and demands of the duty of beneficence aren't relevant. (Whether EA is normative or non-normative, these chapters still represent interesting and useful contributions to the debate around beneficence.)

So some chapters appear to assume that EA makes or should make normative claims about the shape and scope of our duty of beneficence. Other chapters do not. Timmerman argues that EAs should embrace hybridism rather than either actualism (someone ought to do the best of the actions she *would actually* do) or possibilism (someone ought to do the best of the actions she *could* do). This issue clearly arises if EA makes a normative claim (is my obligation to do the most good I *can* do or the most good I *will actually* do?), but Timmerman argues that it arises if EA is a non-normative project too, since someone who has committed to the project of doing good as effectively as possible ought to choose the most effective means to accomplish their goal (and whenever there is an 'ought', even a nonmoral one, questions about actualism and possibilism arise). Collins argues in her chapter that we can sometimes do the most to benefit others by focusing on our duties and responsibilities as members of larger collectives—so political action, which EAs have tended not to focus on, might be more effective than the individual acts of beneficence they tend to emphasize. Sometimes Collins writes in non-normative terms ('individuals who are engaged in the project of benefiting others as much as possible'), but sometimes she writes in normative terms ('the duty to engage in the project of benefiting others as much as possible') (p. 209; p. 208). And the same goes for Chappell's contribution arguing that 'abstract benevolence' is a neglected moral virtue. If EA is merely a non-normative project (if it's just an 'opportunity', not an 'obligation'), then it might be helpful for

me to know whether actualism, possibilism, or hybridism about pragmatic ‘oughts’ is true; whether it would be most efficient for me to act alone or to collaborate; and whether there are any moral virtues that can help me carry out my projects. But if EA makes normative claims, then each of these problems becomes more pressing. If I’m potentially doing something *morally wrong* if I opt for possibilism when hybridism is true, or if I act alone, or if I fail to develop a disposition toward abstract benevolence, then I am morally required to find out what the best interpretation of EA is and to act accordingly.

So we see a theme throughout the volume, especially in the final chapters. Many of the arguments in these chapters apply to non-normative EA—but nearly all of them present us with more pressing issues if EA is telling people what they ought, morally, to do. And some of them (Mogensen’s and Barry and Lawford-Smith’s in particular) don’t seem to apply to EA at all unless it makes normative claims about the duty of beneficence. This picks up, as we have seen, on a divide within the membership of the EA movement and between members of EA and EA leaders. It would make sense for EA to make normative claims, since it is, after all, grounded in philosophical work on the duty of beneficence. On the other hand, there are real benefits to building a truly ecumenical movement that can resist the frequent critique that EA is just act utilitarianism wearing a new hat. Ultimately, I don’t think it’s a flaw in this volume that it doesn’t reach a clear consensus on this issue. Instead, the tensions present in these chapters show that there is much more philosophical work to do on EA. This rewarding volume shows both that EA has made philosophical progress and also that it will continue to be of philosophical interest as we continue to debate how, exactly, we should engage in the project of doing good.<sup>1</sup>

Amy Berg

Rhode Island College

---

<sup>1</sup> I’m grateful to Matt Braich and Theron Pummer for helpful feedback on this review.