The Lessons of Effective Altruism
Jennifer C. Rubenstein*

The Most Good You Can Do, Peter Singer (New Haven, Conn.: Yale University Press, 2015), 232 pp., $16 paper, $25 cloth.

Doing Good Better, William MacAskill (New York: Avery, 2015), 272 pp., $16 paper, $26.95 cloth.

Carol Sue Snowden worked for thirty years as a librarian at the Columbus Metropolitan Library in Columbus, Ohio. She led a quiet, frugal life, spending money mostly on books, which were her passion. When she died, she donated the money she had saved—over $1 million—to the Columbus library and seven local schools.[1] Most of us would look upon this generosity with admiration, but according to a new movement called Effective Altruism (EA), Snowden got it wrong. While she was right to donate her money, she should have instead directed it to an organization that does the *most good* overall.

Roughly three quarters of U.S. households donate to charity, representing a total of almost $265 billion in 2015.[2] Some of these donors are what we might call "passionate philanthropists," like Snowden: they make a deliberate choice to donate a large amount of money to an issue or organization about which they care deeply. Others are what social psychologists call "warm glow" givers. They contribute smaller amounts of money to several different organizations in response to dramatic images of suffering, requests from friends or relatives, and/or to support causes that are close to their hearts. According to EA, both groups are misguided. Individuals should donate to organizations that do the most good— that is, organizations that maximally improve the welfare of sentient beings. While most organizations that Effective Altruists think will do the most good address social problems, such as severe poverty, some Effective Altruists support

more exotic causes, such as slowing human aging, figuring out how to blow up asteroids headed toward Earth, and preventing the "robot apocalypse."[3] Other activities endorsed by Effective Altruists include choosing a career path that will do the most good, not eating meat (especially chicken), and donating a kidney.

Two recent books by philosophers—*The Most Good You Can Do*, by Peter Singer, and *Doing Good Better*, by William MacAskill—aim to introduce EA to an educated lay audience. Despite the more incremental title of MacAskill's book, both authors enthusiastically endorse EA's commitment to doing the most good. As MacAskill writes, the question that Effective Altruists ask themselves is "how can we . . . succeed in having the greatest positive impact we can?" (MacAskill, p. 5; see also MacAskill, p. 12 and Singer, p. vii). Both books show that EA is a better approach to donating than passionate philanthropy and warm-glow giving. However, the distortions that arise from EA's commitment to maximization (doing the *most* good), its narrow focus on improving individual welfare (doing the most *good*), and the disturbing "hidden curriculum" that it teaches (discussed below) suggest the need to move beyond it.

Founding Moments

Effective Altruism had two founding moments. In the mid-2000s, MacAskill and another philosophy graduate student at Oxford, Toby Ord, began researching the cost-effectiveness of charities fighting global poverty. In 2009 they started the EA "meta-charity" Giving What We Can, which encourages people to donate at least 10 percent of their income to the charities with the greatest positive impact. At around the same time, two hedge-fund analysts, Holden Karnofsky and Elie Hassenfeld, independently founded another organization, GiveWell, which focuses on in-depth research to identify the charities that do the most good—a determination that incorporates, but is not limited to, being highly cost-effective (MacAskill, p. 110). Thus, EA came of age simultaneously in Oxford, where it has been especially popular among young analytic philosophers, and in Silicon Valley, where it has been especially popular among young tech entrepreneurs (GiveWell is located in San Francisco). Not surprisingly given this history, the membership of EA organizations has been, and remains, predominantly young, white, well-educated, able-bodied, and (to a lesser degree) male. While Singer is at pains to offer portraits of Effective Altruists who do not fit this description (ch. 3), EA is clearly a product of, and continues to inhabit, these milieus.

Jennifer C. Rubenstein

EA has grown rapidly. At the time of this writing, Giving What We Can has 2,014 members who have collectively donated more than $16 million to "highly effective" charities and have pledged to donate more than $780 million over the course of their careers.[4] GiveWell estimates that it has directed more than $172 million to what it sees as highly effective charities since its founding in 2007.[5] While these numbers are small compared to overall levels of charitable giving, they are significant in absolute terms. They are also only one manifestation of EA's influence, which also includes shifting scholarly and public debates about giving, changing how charities conceptualize and measure their effectiveness, and influencing how some young people choose a career path.

Singer rightly describes himself as a "parent" of the EA movement. In his seminal 1972 essay, "Famine, Affluence, and Morality," he argued that one should donate money to address serious suffering "up to [and perhaps even beyond] the point at which by giving more one would begin to cause serious suffering for oneself and one's dependents."[6] EA retains Singer's emphasis on the demanding moral duty to donate, but adds something new: using reason and empirical evidence to ensure that one's donations do the most good possible. So while in his 1972 essay Singer argued that individuals in wealthy countries should direct their donations to help victims affected by the "emergency" of "constant poverty, a cyclone, and a civil war" in East Bengal, in *The Most Good You Can Do* he clearly supports an argument that MacAskill makes explicit—that individuals should generally *not* provide emergency relief in response to high-profile disasters because it tends to do less good than other interventions (MacAskill, ch. 4).

Singer and MacAskill are both utilitarians. While they seek to build a big tent and persuade people who do not share their utilitarian commitments, the shadow of utilitarianism looms large over both books. Both authors view all human individuals as having equal and great moral worth, a commitment shared by many moral theories. However, their focus on maximally increasing individual welfare as the most appropriate way to enact this commitment is most closely associated with utilitarianism. EA can therefore trace its ancestry not only to Singer but also to earlier utilitarian philosophers, such as John Stuart Mill and Jeremy Bentham, who argued for bringing about "the greatest amount of good for the greatest number."[7]

A Philosophy and a Movement

While it has a philosophical lineage, EA is not only a set of philosophical ideas. It is, as Singer emphasizes, "a philosophy *and* a movement" (p. 5, emphasis added). It consists of both abstract principles and efforts to enact those principles, including EA meta-charities, EA-influenced charities, EA blogs, discussion groups, chapters, TED Talks, conferences, scholarly papers, and websites offering career guidance. This proliferation raises the sticky question of how books that seek to popularize EA, such as those under consideration here, should be evaluated.[8] Should we focus only on the core ideas of EA, for example that individuals should do as much good as possible? Or should we also include what happens when these ideas are put into practice?

I propose that, as a philosophy and movement that focuses on consequences, EA should be evaluated in light of all of its overall effects, even those that run counter to its stated intentions. For example, EA meta-charities are in principle open to endorsing any organization that can be shown to do the most good. In practice, however, they systematically prioritize organizations that provide public health interventions and cash transfers over those that engage in political advocacy, because the effects of the former are easier to measure and compare than those of the latter (MacAskill, pp. 37–38). While many within EA are attempting to address this issue, it cannot be dismissed on the grounds that it is unintended. The regular but unintended effects of EA are features—not minor bugs—of the system.

The Argument for Effective Altruism

Singer and MacAskill make at least three important arguments in favor of EA. These arguments are primarily empirical, but they have powerful normative implications. First, there are huge differences in the good that can be done by different interventions (MacAskill, ch. 4). Second, these differences are often measurable. While these measurements are typically quite rough, the differences among various kinds of interventions are sometimes so great that they can still be captured using highly inexact methods. For example, a donation of $50,000 can be used to train and provide one guide dog for one blind person in the United States, or it can be used to fund surgery to *prevent* blindness (from trachoma) for *five hundred* people in a poor country (MacAskill, p. 61).[9] Singer and MacAskill's third important argument is that most people in wealthy countries

are very well-off in absolute terms, and vastly better off than many people in poor countries (MacAskill, ch. 1). As MacAskill points out, a worker in the United States who earns $52,000 a year is still in the top 1 percent of earners globally (p. 18).

Together, these three arguments suggest that most adults in wealthy countries can, without immiserating themselves, donate enough to make a meaningful difference in the lives of many poor people. Even if MacAskill's estimation that it costs $3,400 to save a life in the developing world (p. 54) is optimistic, it casts passionate philanthropy and warm-glow giving (as well as the refusal to donate at all) in a new light. If a young adult with trachoma, facing a lifetime of blindness, asks you why you are funding one five-hundredth of a guide dog rather than a surgery to enable her to see, what would you say? Singer and MacAskill do not show that there are no good answers to this question. But they do show, powerfully, that it is a question we should ask.

More prosaically, the idea that donors should do the most good generates insights about everyday donating practices that are relevant for anyone who cares about the effects of their donation, not only Effective Altruists. Donors are frequently told that, when deciding where to donate, they should consider how much of an organization's budget is devoted to overhead versus program costs. EA tells us that this figure is virtually meaningless; what matters is the amount of good that an organization accomplishes (and the harm it inflicts) per dollar spent, regardless of whether that dollar is categorized as overhead or program cost (MacAskill, ch. 7). Likewise, while donors frequently prioritize causes in which their donation will address all or almost all of a problem, EA suggests that we should focus instead on the absolute size of our contribution. As MacAskill writes, "It's not the size of the bucket that matters, but the size of the drop" (p. 25).

Criticisms of Effective Altruism

Despite the power and usefulness of its arguments, EA has no shortage of detractors. I turn now to discussing three lines of criticism. The first is a common accusation that, I suggest, EA has the resources to parry; the latter two pose a more serious challenge.[10]

Intimate Donating
One common criticism of EA involves what we might call "intimate donating," that is, donating that is intended not only to support the aims of the recipient

organization but also, and sometimes primarily, to express or do something within an intimate relationship. For example, when I donate to my niece's school dance-a-thon for the American Cancer Society or participate in the Ice Bucket Challenge at the behest of an old friend, a big part of what I am doing is expressing my love for my niece and connecting with my friend. What can EA say to donors who understand and accept that they are not doing as much good as possible, but take themselves to be doing something else that is also morally valuable?

EA's main response to this question has been developmentalist. At the height of the Ice Bucket Challenge craze, MacAskill urged people to conceive of their participation in it not as "'doing [your] bit' but instead as taking one small step towards making altruism a part of your identity."[11] That is, while participating in the challenge does not do the most good (because the recipient organization, the ALS Association, is not one of the most effective charities), it could be part of an individual's process of *moving toward* doing the most good. Throughout their books, both Singer and MacAskill describe how aspiring, or even tentatively aspiring, Effective Altruists can move gradually toward a more EA-consistent lifestyle, and they encourage their readers to follow suit.

This developmentalist response seems to imply that full compliance with EA principles is incompatible with intimate donating. This is not the case, however. As Singer notes, "Typical effective altruists leave themselves time and resources to relax and do what they want" (p. 8). So as long as donations made via the Ice Bucket Challenge are classified as "luxury spending," and therefore do not displace donations (or other activities) that *do* do the most good, there is no reason to object to them (p. 45).[12] The same point holds for other donations that do not do the most good.

A critic might rejoin that there are important differences in kind between participating in the Ice Bucket Challenge and, say, buying oneself an expensive stereo. Effective Altruists do not need to deny this, however. Their argument suggests that the *bolded* boundary on our moral maps ought to be drawn in a different place than we usually draw it, not that other kinds of boundaries do not exist. They aim to shift attention away from the distinction between all kinds of do-gooding (including doing the most good) and self-interest, and toward the distinction between doing the most good and everything else (including both self-interest and other kinds of do-gooding). This does not commit them to denying the existence of salient distinctions within the latter category.

Jennifer C. Rubenstein

If all donations did about the same amount of good (and no harm), there would be no reason *not* to deploy one's charity dollars to, in effect, do double duty by both expressing love or enacting friendship while also contributing to the broader social good. However, EA tells us that because there are vast differences in the good that different kinds of donations do, this kill-two-birds-with-one-stone strategy involves a steep moral cost. The more general point that I take from this is not that leading a life fully devoted to nurturing deep friendships—or to playing music, training guide dogs, or undertaking basic scientific research—is wrongheaded, as some Effective Altruists imply, but rather that such a life sacrifices what Effective Altruists argue persuasively is a distinctive and important value: doing as much good as you can.

Effective Altruism Is Anti-Political

A second and more serious criticism of EA is that it offers technical and economic solutions to what is, in fact, the *political* problem of severe poverty.[13] Among other things, critics allege that EA is biased against advocacy for political change. As evidence, critics typically point out that, as noted above, most organizations identified as "most effective" by EA meta-charities focus on public health and cash transfers, not political advocacy. One explanation for this (which Effective Altruists readily acknowledge) involves measurement bias: because EA aims to do the *most* good, it needs to measure and compare the good that different interventions accomplish (MacAskill, pp. 38–39). This objective helps to drive EA's singular focus on doing good and its reticence to consider other values, which are harder to measure and commensurate. It also gives EA a reason to conceptualize "doing good" as increasing individual welfare: because individual welfare can be measured (sort of) in "quality-adjusted life years" or similar units, the effects of initiatives that aim to increase individual welfare can be compared—for example by using randomized controlled trials (MacAskill, pp. 39–40). Because public health interventions and cash transfers are more amenable to these approaches than political advocacy, EA meta-charities have tended to focus on them.

Why is EA's relative inattention to political advocacy objectionable? Critics offer two distinct reasons. One is that political advocacy addresses "root causes," and so is more cost-effective and sustainable over the long term than the "Band-Aid" solutions that EA supports, and thus is better than the approaches that EA endorses by its own criteria. However, I think EA can respond to this version of the objection fairly easily. Given choices between sure bets that do less

good and long shots that do more good, and between projects that do less good but have effects that are easier to measure and projects that probably do more good but have effects that are harder to measure, EA can choose the latter options. That is, it can put relatively more weight on the former of two values that it already endorses: doing the most good and certainty about doing good.

Another objection to EA's inattention to political advocacy is much more challenging. According to this objection, the problem is not simply measurement bias; it is a more thoroughgoing anti-political sensibility. Indeed, while both Singer and MacAskill explicitly discuss political advocacy campaigns in their books (Singer, pp. 157–64; MacAskill, pp. 94, 174–75), both seem to see the world fundamentally through the lens of individual needs and welfare, rather than power, inequality, injustice, exploitation, and oppression. This sensibility seems to contribute to, and not (only) result from, a judgment that interventions aimed at addressing the latter sorts of issues are risky and unlikely to succeed. For example, Singer writes that "it isn't clear that making the rich richer without making the poor poorer has bad consequences, overall. It increases the ability of the rich to help the poor, and some of the world's richest people, including Bill Gates and Warren Buffett, have done precisely that" (p. 50; see also p. 113). But the salient question about making the rich richer is not, as Singer implies, whether things would be better under global communism (p. 50); it is, rather, about whether and how extreme economic inequality contributes to political inequality and domination.

Altering this anti-political sensibility and its downstream effects would, I think, require a transformational shift in EA. One way to enact this transformation, which I discuss in more detail below, would be for EA to explicitly incorporate other values in addition to increasing individual welfare, such as justice, fairness, and inclusion. (Singer notes that some Effective Altruists already care about these other values as ends in themselves, but he himself seems not to [p. 9]). This more pluralistic approach is philosophically coherent and provides a powerful basis for critiquing passionate philanthropy and warm-glow giving. The bigger challenge involves its implications for EA as a movement. Among other things, EA organizations such as GiveWell and Giving What We Can currently offer (a) concrete, easy-to-follow guidance to donors about where to donate and (b) the opportunity to read and/or engage in discussions that draw heavily on the tools of economics and analytic philosophy. Shifting the central objective of EA to "doing a lot of

Jennifer C. Rubenstein

good and attending to other values" would require a more extensive rethinking of both offerings.

To be clear, adopting a more political sensibility is not the same as designating more organizations that engage in political advocacy as "most effective." Even when it involves empowering marginalized groups, political advocacy by outsiders can amount to democratically illegitimate meddling.[14] It can also distort the agendas of domestic organizations or individuals that have a better understanding of local conditions than the external organizations that fund them.[15] Even advocacy aimed at reforming the harmful policies and practices of a donor's own government, which some critics of EA have endorsed,[16] is not without moral danger. For example, Section 1502 of the Dodd-Frank Act requires U.S. companies to show that minerals they import from conflict-prone areas in the Democratic Republic of the Congo are not "conflict minerals." Yet despite good intentions, activists from both the DRC and the United States have argued that not only does this do little to stem the conflict, it also makes life much worse for the miners and their families.[17] All of which is to say that the question should not be "How can EA meta-charities tweak their criteria so they end up endorsing at least some organizations that do political advocacy?" but rather "How can EA as a whole ensure that its *analysis* is attentive to political dynamics—even if it turns out that the best role for outsiders is addressing symptoms rather than underlying causes?"

Effective Altruism's "Hidden Curriculum"

My own deepest worries about EA pertain not to the interventions that it supports, but rather to the self-understandings and sensibilities that it cultivates in its members via its "hidden curriculum." A hidden curriculum is the set of "unwritten, unofficial, and often unintended lessons, values, and perspectives" that are taught alongside an explicit official curriculum.[18] EA teaches a hidden curriculum to its members through its arguments and images, the emotions it evokes, and the relationships that it fosters and discourages.

As I discuss below, EA's hidden curriculum includes at least four lessons: (1) Effective Altruists are heroic rescuers; (2) doing good is largely an individualistic project; (3) doing the most good does not require listening to those affected by the issues one is trying to address; and (4) anger is not an appropriate response to severe poverty. I am not suggesting that Effective Altruists consciously believe these lessons or teach them explicitly; nor can I offer here evidence about the extent to which they have been taken up. But I do think that EA conveys them.

Moreover, these lessons are not mere window-dressing; they serve important functions in helping EA attract and retain members. Let us look at them more closely.

Lesson 1: Effective Altruists are heroic rescuers. EA's hidden curriculum teaches people in wealthy countries to think of themselves as heroes and rescuers. In the opening paragraph of his book, Singer recounts a realization that one of his students had that led him to adopt an EA lifestyle: "Suppose you see a burning building and you run through the flames and kick a door open, and let one hundred people out. That would be the greatest moment in your life. And I could do as much good as that!" (p. 3). Singer tries to motivate his readers to become Effective Altruists by encouraging them to adopt this same outlook—to see that they, too, can easily rescue many people. MacAskill does the same thing: "You pass a burning building, kick the door down, rush through the smoke and flames, and drag a young child to safety. . . . If you saved several people's lives [in this way] . . . you'd be a hero." And then he adds, "But we can do *far* more than that" (pp. 53–54).

This kind of high-drama emergency rescue scenario is powerfully motivating (which might be why Singer consistently invokes such scenarios in his work).[19] However, it might also encourage aspiring Effective Altruists to think of themselves as rescuers, and the people they wish to assist as helpless victims more generally.[20] This conception of "self and other" can have several negative and distorting effects. It can make it harder for the self-described rescuer to notice the ways in which she has contributed to and/or benefited from the problems she seeks to address, and it can lead her to discount the insights of the "victims." The desire to have *personally* made a quantifiable difference can contribute to a bias in favor of interventions, such as public health interventions, that can be plausibly parceled out in ways that give credit to individual donors (for example, "Your donation will save ten lives!").

An obvious response to this criticism is that EA should simply use different tactics to motivate its members. While I agree that this is a good idea, a central strength of EA is that it eschews several other motivational strategies that are equally or more objectionable—for example, showing images of pitiful-looking "identifiable victims." While some form of political solidarity is likely preferable, the fact that EA focuses on doing the most good means that Effective Altruists have to be ready to cut bait if a particular organization or cause ceases to do

Jennifer C. Rubenstein

the most good, which significantly complicates political solidarity as a motivational strategy.

Lesson 2: Doing good is a largely individualistic project. A second lesson in EA's hidden curriculum is that doing good is an individualistic project. Of course, heroes are lone figures, so this lesson overlaps with the previous one. But here I emphasize that EA assimilates the role of donor to two other individualistic roles. One is the savvy customer: MacAskill notes that donors in rich countries can provide "100x" more benefit by donating to the most cost-effective charities than to other ones. This, MacAskill exclaims, "might be the most amazing deal you'll see in your life" (p. 23). The other role is the virtuous self-improver: This idea—that Effective Altruists ought to be engaged in a continual process of self-conscious self-improvement—emerges out of the developmental story that both MacAskill and Singer tell, discussed above. As Singer puts it, Effective Altruists are people who "like to challenge themselves, to do a little better this year than last year" (p. viii).

All three of these aspirational figures—the hero, the savvy consumer, and the virtuous self-improver—direct the donor's focus back toward herself, and away from collective practices of joint action with others. I worry that if doing good is conceived of as not only the most important project but also as akin to heroism, bargain-hunting, and self-improvement, then more solidaristic modes of action—which have historically undergirded many projects that have significantly improved the human condition, such as struggles for civil rights—will cease to be easily imaginable.

Lesson 3: Efforts to do the most good do not require listening to or working alongside those affected by the issues one is addressing. While EA is individualistic in some respects, it is also a "community" (MacAskill, p. 12). The EA meta-charity Giving What We Can intentionally fosters community among its members both to motivate them and to reduce the social burdens of an EA lifestyle (such as not going out to expensive restaurants). Yet unlike many other identity-based communities, the EA community systematically excludes the very people who are its central objects of concern. While alleviating severe poverty is a primary aim of EA, poor people generally do not have enough money or time to become Effective Altruists. Likewise, while there are exceptions, poor people and activists in poor countries often find themselves already enmeshed in the particular issues that

affect their communities, such that one of EA's central questions—Which of all of the issues and causes in the world should I support in order to do the most good? —is not salient for them.

As a result, when members of the EA community discuss issues among themselves, they are not talking to poor people or to activists who work closely with them (as far as I have been able to discern). Nor does there seem to be any effort to put the EA community in contact with activists, civil society groups, or NGOs based in poor countries. I tend to think that this is objectionable in itself. But in addition, an important limitation of the methods that EA uses to evaluate different interventions, such as randomized control trials, is that they do a poor job of capturing long-term, diffuse, and subtle political and social effects—for example, on intergroup social relationships, state capacity, and civil society mobilization. Engagement with those affected and others with more direct experience of the issues that EA seeks to address would likely be extremely valuable in helping them account for these effects. While Effective Altruists have increasingly noticed the limitations of their formal models and the need to rely more on "intuition" and "reality checks,"[21] they do not seem to have structures in place to take up experiential, anthropological, or humanistic modes of knowledge (a notable exception is the Open Philanthropy Project's work on the history of philanthropy).[22]

In addition to potentially undermining EA's efforts to achieve its own goals, this sort of insularity can contribute to a more generalized myopia. For example, at a recent global conference on Effective Altruism an event organizer declared in an introductory session that he "really [does] believe that Effective Altruism could be the last social movement we ever need."[23] Given that EA has virtually no interest in, or capacity for, addressing issues of injustice, exploitation, and oppression, one wonders whether more time spent listening to people facing those issues might have led the speaker to a different conclusion.

Lesson 4: Anger is not an appropriate response to the issues that EA addresses. While Effective Altruists are often accused of being cold rationalists, and while they do prioritize reason over emotion, they do in fact recognize a role for emotions (Singer, p. 83; see also pp. 85, 88, 92–93, 171, and MacAskill, pp. 5, 11). What is the emotional content of EA? Recall Singer's student, noted above, who declared that saving lives "would be the greatest moment" in his life (see also p. 133). Singer also describes another Effective Altruist who gives "cheerfully" (p. 29). Singer adds that "effective altruists don't see a lot of point in feeling guilty.

Jennifer C. Rubenstein

They prefer to focus on the good they are doing. Some of them are content to know they are doing something significant to make the world a better place" (p. viii). Finally, Singer emphasizes that one can be an Effective Altruist without feeling grim and depressed about living on less money than one would otherwise (pp. 23–37, ch. 9). (Anecdotally, when I submitted a blog post to the Giving What We Can Blog, I received instructions for authors explaining that "because we are welcoming, we use photos of our members: smiley, happy people.")

There thus seems to be little room for emotions such as anger, guilt, and self-doubt in EA. While Singer and MacAskill do not say that these emotions are bad or illegitimate, they strongly emphasize the value and appropriateness of their opposites (MacAskill, p. 32). Of course, it is true that, as Martha Nussbaum argues, anger can direct energy to retribution rather than constructive efforts for change.[24] However, as Nussbaum also writes, "Anger is sometimes justified and right. It is an appropriate response to injustice and serious wrongdoing."[25] On a cognitive view of emotions, being angry about an injustice can be partly constitutive of recognizing it as such. Likewise, guilt can be a part of acknowledging that one has wrongly contributed to suffering, and self-doubt can reflect humility and an understanding that one does not have all of the answers.

Perhaps, then, EA could create a bit more space for these emotions? Among other things, recognizing the appropriateness of some kinds of anger would be consistent with a more politicized analysis of the issues that EA addresses: what looks like a regrettable misfortune when we focus only on individual welfare (a child gets malaria) becomes recognizable as an injustice (to which anger is an appropriate response) when we look at the underlying structures that contribute to it (an externally enabled corrupt ruling elite that drains resources from government public health budgets). Openness to negative emotions may be challenging, however, insofar as members of EA are attracted to its cheery emotional register, and anger at injustice can be difficult to square with willingness to cut bait if an intervention does not seem to be doing the most good.

EA's hidden curriculum is deeply troubling, in part because it is not window-dressing: it does important work, especially motivational work. But this should not cause us to lose sight of EA's central insight, described above: we should be thinking rigorously, comparatively, and collectively about how individuals can use the money in their possession to improve the world, including (but not only) by increasing individual welfare.

Another Way Forward?

What might it look like for a philosophy and social movement to grow up around *this* insight, without EA's narrow focus on either maximization (doing the *most* good) or welfare (doing the most *good*)? Can we imagine an organization, very roughly analogous to GiveWell, that would offer guidance to individual donors who reject warm-glow giving and passionate philanthropy, and are looking instead for guidance on donating that is rigorous, empirically grounded, comparative, and attentive to consequences but that also acknowledges and incorporates a range of hard-to-measure substantive and procedural values, including justice, equality, fairness, and empowerment of those directly affected?

Critics of charity will object that in a just world private individuals would not be able to accumulate enough money to make voluntary donating a significant driver of social change, and in the current (unjust) world the wealthy are unlikely to support reforms that would limit their ability to accumulate.[26] I think this objection is overstated, both because an effective social practice of donating does not require the excessive accumulation that we associate with such megadonors as Bill Gates and Warren Buffett, and because individuals often donate to support political reforms that will reduce their own ability to accumulate.[27]

Any effort to create a more pluralistic, less maximizing GiveWell would of course learn from and build on the work of many existing organizations, such as Solidaire, which funds social change organizations in the United States;[28] Thousand Currents (formerly IDEX), which gives grants to grassroots organizations in poor countries and accepts donations from the general public;[29] and Shack/Slum Dwellers International and WIEGO, both of which empower and give voice to very poor people.[30] Such an effort would also build on the insights, experiences, and arguments of EA, which bring into view the following crucial questions that would have to be answered in order to ascertain whether the sort of GiveWell alternative that I have been describing would be possible or beneficial:

- How can we create criteria for evaluating organizations, interventions, or causes that (1) make comparison possible, but (2) are responsive to the plural goals of the organizations being evaluated, the people they aim to assist, and the donors who will use these criteria to inform their donating practices?

Jennifer C. Rubenstein

- What are the precise mechanisms whereby donations from individuals in wealthy countries can promote not only individual welfare but also justice, equality, and/or other values? What are the best sources for learning about these mechanisms?
- How can organizations based in the Global South accept donations from individuals in the Global North without compromising their actual and perceived democratic legitimacy?[31]
- Given that several of the negative lessons in EA's hidden curriculum emerge from its efforts to motivate its members, would this alternative organization simply serve as a resource for donors who are already motivated? How would it balance solidarity and partnership with the willingness to shift gears?

Singer and MacAskill have done an outstanding job of presenting EA to a wide audience. For those of us who see the promise in EA, but also recognize its limitations, their books are best seen as contributions to a much larger conversation.

NOTES

1. "Librarian Quietly Saved $1 Million For Gift Back To Library," *Huffington Post*, September 12, 2013, www.huffingtonpost.com/2013/09/12/library-donation_n_3915654.html.
2. Amount donated by individuals according to Charity Navigator, www.charitynavigator.org/index.cfm?bay=content.view&cpid=42 (accessed September 22, 2016). Estimations of the percentage of the U.S. population that gives range from 67 percent of households (Philanthropy Roundtable, www.philanthropyroundtable.org/almanac/statistics/) to 83 percent of adults (Gallup, "Most Americans Practice Charitable Giving, Volunteerism," December 13, 2013, www.gallup.com/poll/166250/americans-practice-charitable-giving-volunteerism.aspx). Both sites accessed September 22, 2016.
3. The term "robot apocalypse" is from Amia Srinivasan, "Stop the Robot Apocalypse," *London Review of Books* 37, no. 18 (2015) pp. 3–6. See also the Open Philanthropy Project, "Global Catastrophic Risks," February 2014, www.openphilanthropy.org/research/cause-reports/global-catastrophic-risks/global-catastrophic-risks (on asteroids); and Zander Redwood, "Living to 1,000: An Interview with Aubrey de Grey," 80,000 Hours, April 12, 2012, 80000hours.org/2012/04/living-to-1000-an-interview-with-aubrey-de-grey/ (on aging).
4. See the Giving What We Can website, "Our History," www.givingwhatwecan.org/about-us/history/.
5. See the GiveWell website, "GiveWell's Impact," www.GiveWell.org/about/impact.
6. Peter Singer, "Famine, Affluence, and Morality," *Philosophy and Public Affairs* 1, no. 3 (1972), pp. 229–43.
7. Julia Driver, "The History of Utilitarianism," in Edward N. Zalta, ed., *Stanford Encyclopedia of Philosophy*, Winter 2014 Edition, plato.stanford.edu/archives/win2014/entries/utilitarianism-history/.
8. It also means that the positions of EA organizations, and discussions of EA on blogs and websites, have already moved beyond some of the ideas in these books. For example, while MacAskill describes "earning to give" (choosing a high-earning job so that you can give away more money) as a "solid bet" (pp. 162–63), the EA organization that proposed it (80,000 Hours) has pedaled back its support for this approach. See William MacAskill, "80,000 Hours Thinks that Only a Small Proportion of People Should Earn to Give Long Term," 80,000 Hours, July 6, 2015, 80000hours.org/2015/07/80000-hours-thinks-that-only-a-small-proportion-of-people-should-earn-to-give-long-term/. While this review responds primarily to the version of EA presented by MacAskill and Singer, readers who want an up-to-date understanding of the positions and debates circulating within EA should consult blogs and websites of EA and EA-affiliated organizations, such as GivingWhatWeCan.org and GiveWell.org.

9 Some empirical scholars doubt the veracity of these two arguments. See Angus Deaton, Responding to Peter Singer in the *Boston Review* online, July 1, 2015, bostonreview.net/forum/logic-effective-altruism/angus-deaton-response-effective-altruism.

10 The best critical discussions of EA (some more critical than others) include those on Amia Srinivasan, in the "Thread for Discussing Critical Review of Doing Good Better in the London Review of Books," September 21, 2015, effective-altruism.com/ea/nw/thread_for_discussion_of_critical_review_of_doing); Iason Gabriel, "Effective Altruism and Its Critics," *Journal of Applied Philosophy* 33, no. 3 (2016); Dylan Matthews, "I Spent a Weekend at Google Talking with Nerds About Charity. I Came Away . . . Worried," *Vox*, August 10, 2015; Derek Thompson, "The Greatest Good," *Atlantic*, June 15, 2015; and Robert Simpson, "Moral Renegades," *New Rambler Review of Books*, July 12, 2016, newramblerreview.com/book-reviews/philosophy/moral-renegades.

11 William MacAskill, "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/.

12 The term "luxury spending" is not Singer's, but he cites it approvingly.

13 Deaton, Responding to Peter Singer, July 1, 2015.

14 Ted Lechterman, "The (Effective) Altruist's Dilemma," unpublished manuscript (August 2016).

15 For example, Nandini Deo and Duncan McDuie-Ra describe how, in order to please external funders, NGOs in India prioritized interventions aimed directly at ending violence against women, rather than addressing the underlying cause of such violence, which was women's lack of access to employment. Deo and McDuie-Ra, *The Politics of Collective Advocacy in India: Tool and Traps* (Sterling, Va.: Kumarian Press, 2011), ch. 2.

16 Deaton, Responding to Peter Singer, July 1, 2015.

17 Laura Seay, "What's Wrong with Dodd-Frank 1502? Conflict Minerals, Civilian Livelihoods, and the Unintended Consequences of Western Advocacy," Center for Global Development Working Paper 284, January 2012, www.cgdev.org/publication/what%E2%80%99s-wrong-dodd-frank-1502-conflict-minerals-civilian-livelihoods-and-unintended.

18 See the website of the Glossary of Education Reform, "Hidden Curriculum," edglossary.org/hidden-curriculum/.

19 Peter Singer, *The Life You Can Save* (New York: Random House, 2010).

20 Gabriel, "Effective Altruism and Its Critics."

21 The GiveWell Blog, "Why We Can't Take Expected Value Estimates Literally (Even When They're Unbiased)," blog.givewell.org/2011/08/18/why-we-cant-take-expected-value-estimates-literally-even-when-theyre-unbiased/, and "Some Considerations against More Investment in Cost-Effectiveness Estimates," blog.givewell.org/2011/11/04/some-considerations-against-more-investment-in-cost-effectiveness-estimates/. Both posts updated July 25, 2016.

22 Open Philanthropy Project, "History of Philanthropy," www.openphilanthropy.org/research/history-of-philanthropy?expand_menu=1.

23 Matthews, "I Spent a Weekend at Google Talking with Nerds About Charity."

24 Martha Nussbaum, *Anger and Forgiveness: Resentment, Generosity, Justice* (New York: Oxford University Press, 2016).

25 Martha Nussbaum, *Political Emotions: Why Love Matters for Justice* (Cambridge, Mass.: Belknap Press, 2015), p. 394.

26 Mathew Snow, "Against Charity," *Jacobin*, August 25, 2015, www.jacobinmag.com/2015/08/peter-singer-charity-effective-altruism/.

27 Solidaire, solidairenetwork.org/.

28 Ibid.

29 "Thousand Currents: Exchanging Grassroots Brilliance," as described on the homepage of the International Development Exchange (IDEX), www.idex.org/.

30 Shack/Slum Dwellers International (SDI), "About SDI," sdinet.org/about-us/what-we-do/; and Women in Informal Employment: Globalizing and Organizing (WIEGO), wiego.org/.

31 Deo and McDuie-Ra, *The Politics of Collective Advocacy in India.*

Jennifer C. Rubenstein