

CHAPTER 29

EFFECTIVE ALTRUISM

A Consequentialist Case Study

JUDITH LICHTENBERG

1.

EFFECTIVE altruism (EA) emerged over the last decade out of the ideas and writings of two young Oxford philosophers, Toby Ord and William MacAskill. In 2009, Ord founded Giving What We Can, an organization whose members pledge to donate 10 percent of their income to “whichever organisations can most effectively use it to improve the lives of others.” As of June 2019, it had over 4,000 members who have pledged over \$126,000,000 to charity.¹ MacAskill and Benjamin Todd started 80,000 Hours (the average number of hours a person works in her lifetime) in 2011 to help people figure out how to best use those hours “to solve the world’s most pressing problems.” They report that “More than 3000 people have told us that, due to engaging with us, they have significantly changed their career plans and now expect to have a larger social impact as a result.”²

The movement has its roots in the works of Peter Singer, beginning with his classic 1972 essay “Famine, Affluence, and Morality,” a staple of introductory ethics courses ever since. Writing in the wake of a Bangladesh famine, Singer reset the course of contemporary moral philosophy—and certainly the teaching of it—with his deceptively simple premise that “if it is in our power to prevent something bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it.”³ With global poverty as his focus, Singer concludes without much ado that “we ought to give until we reach the level . . . at which, by giving more, I would cause as much suffering to myself or

¹ Giving What We Can, <https://www.givingwhatwecan.org/about-us/>. In 2016, Giving What We Can was incorporated into its parent organization, the Centre for Effective Altruism in Oxford.

² 80,000 Hours, at <https://80000hours.org/>. This organization is also sponsored by the Centre for Effective Altruism.

³ Peter Singer, “Famine, Affluence, and Morality,” *Philosophy & Public Affairs* 1 (1972), 231.

my dependents as I would relieve by my gift.” In so doing, Singer, an avowed utilitarian, married his philosophical theory to the so-called demandingness problem. Ever since, moral philosophers have been asking whether it is reasonable or morally right to expect people to make such significant sacrifices to alleviate the suffering of others.

Singer’s 2015 book *The Most Good You Can Do: How Effective Altruism Is Changing Ideas about Living Ethically* puts the point more clearly and explicitly. In contrast to the original, negative formulation in terms of “preventing something bad from happening,” it expresses the point positively—“Effective altruism is based on a very simple idea: we should do the most good we can.”⁴ In his 2015 book *Doing Good Better: How Effective Altruism Can Help You Make a Difference*, MacAskill similarly asserts that EA is about making the biggest difference you can. “Of all the ways in which we could make the world a better place, which will do the most good?”⁵

Singer’s and MacAskill’s books, and other products of the EA movement, have undoubtedly generated some very good results. For one thing, they help rebut the scores of books and articles published over the last several decades arguing that international “aid” is at best ineffective and at worst counterproductive. These books have names like *The White Man’s Burden*, *Dead Aid*, *The Road to Hell*, *Lords of Poverty*, *Famine Crimes*, and *The Dark Sides of Virtue*.⁶ Aid’s critics do not really deny that it’s possible for affluent westerners to improve the lives of the world’s poorest people—improving health outcomes is the most obvious example—but that fact often gets obscured by the titles and headlines. Effective altruists help set this record straight. They show how without enormous sacrifices people can greatly improve the odds that their donations will make substantial improvements in human well-being.

But EA, I shall argue, rests on problematic assumptions. Are effective altruists right that one should always do the most good one can? Even if the answer is no, is doing more good always morally better than doing less? Before tackling these questions I consider a prior one: does EA necessarily presuppose consequentialism?

⁴ Peter Singer, *The Most Good You Can Do: How Effective Altruism Is Changing Ideas about Living Ethically* (New Haven, CT: Yale University Press, 2015), vii.

⁵ William MacAskill, *Doing Good Better: How Effective Altruism Can Help You Make a Difference* (New York: Gotham Books, 2015), 32. In that sense the title of his book is misleading. A person might well do good *better* than she had but still not make the biggest difference she could. I discuss this ambiguity below.

⁶ See William Easterly, *The White Man’s Burden: Why the West’s Efforts to Aid the Rest Have Done So Much Ill and So Little Good* (New York: Penguin Press, 2006); Dambisa Moyo, *Dead Aid: Why Aid Is Not Working and How There Is a Better Way for Africa* (New York: Farrar Straus & Giroux, 2009); Michael Maren, *The Road to Hell: The Ravaging Effects of Foreign Aid and International Charity* (New York: Free Press, 1997); Graham Hancock, *Lords of Poverty: The Power, Prestige, and Corruption of the International Aid Business* (New York: Atlantic Monthly Press, 1989); Alex de Waal, *Famine Crimes: Politics and the Disaster Relief Industry in Africa* (London: Africa Rights and the International African Institute, 1997); David Kennedy, *The Dark Sides of Virtue: Reassessing International Humanitarianism* (Princeton, NJ: Princeton University Press, 2005). For further discussion of these critiques, see Judith Lichtenberg, *Distant Strangers: Ethics, Psychology, and Global Poverty* (Cambridge: Cambridge University Press, 2014), chap. 8.

2.

I take consequentialism to be the view that the rightness or wrongness of acts depends solely on their consequences—their tendency to increase the good or decrease the bad, however good and bad are understood. Utilitarianism, the classic form of consequentialism, identifies the good with pleasure or happiness or well-being and the bad with its opposite. Utilitarianism has declined in popularity over the last few decades and has been replaced among consequentialists by more pluralistic accounts of the good.

Singer is an avowed consequentialist (a utilitarian, even), and many other prominent advocates probably are as well. Whether consequentialism is an essential feature of EA depends what one takes the latter's core ideas to be. If the point is simply to do more to address the world's greatest problems, such as global poverty, and to do it more effectively, the answer is no.⁷ Over the last several thousand years, many religious, ethical, and political traditions have advocated remedying poverty and inequality. The Bible urges us to "Sell what you have and give the money to the poor."⁸ Augustine asserts that "The superfluity of the rich is necessary to the poor. If you hold onto superfluous items, then you are keeping what belongs to someone else."⁹ Liberals, socialists, Marxists, communists, and others have, on moral grounds, long called for a more equitable distribution of wealth that allows the least well off to live decently (if not equally). Even economists, invoking the principle of diminishing marginal utility, find good reasons to defend redistribution from richer to poorer (*ceteris paribus*). So the idea that those who are able to ought to do more, even much more, to alleviate poverty, and that we should not waste our efforts, are hardly unique insights of EA and do not entail consequentialism.

How, then, does EA differ from these age-old injunctions, if it does? The suspicion is that it presents us with an extremely demanding command to maximize the good (or minimize the bad), as its godfather, Peter Singer, has urged. The "demandingness problem" has beset contemporary ethics, and the maximizing version seems to be the default interpretation of consequentialism today.

In a brief essay sympathetic to EA, Jeff McMahan denies that the view presupposes consequentialism. Referring to Singer's argument in "Famine, Affluence, and Morality," McMahan says it "appealed in the first instance to a single widely held moral intuition and argued that consistency required those who accepted the intuition to give most of their wealth to the relief of extreme poverty."¹⁰ Likewise for Peter Unger's argument for a similar conclusion in *Living High and Letting Die*, which, McMahan says, explicitly disavowed commitment to any particular moral theory. "His aim was to demonstrate that a

⁷ I shall focus here on the alleviation of global poverty, one of EA's main concerns. It's not the only one, however, and not everything I say here will apply to all of the movement's targets.

⁸ Luke 12:33.

⁹ Augustine, *Exposition of the Psalms 121–150*, vol. III/20 (Hyde Park, NY: New City Press, 2004).

¹⁰ "Philosophical Critiques of Effective Altruism," <http://jeffersonmcmahan.com/wp-content/uploads/2012/11/Philosophical-Critiques-of-Effective-Altruism-refs-in-text.pdf>, 1.

view of the sort that now informs the work of effective altruists is implicit in values and convictions we already have.”¹¹

But do Singer’s and Unger’s conclusions follow from “a single widely held moral intuition” or from “values and convictions we already have”? Equally plausible is that once we understand what these philosophers think these values and intuitions imply, what appears to be a widely held moral intuition (such as “if it is in our power to prevent something bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it”) is one we are ready to discard. We see that we misunderstood this “intuition” and do not in fact accept it. That, at least, is what I take away from my decades of teaching Singer’s article to thousands of undergraduates. To assert that our assent to Singer’s principle on first and casual reading amounts to a widely held intuition or deep-seated value is simply wrong.¹² On a pragmatic understanding of meaning, we do not accept this principle if we reject what seem to be its obvious implications.

So do effective altruists insist that we do the most good we can, as Singer’s title implies? Or should we merely “do good better,” as MacAskill’s title suggests? If the latter, is that anything more than common sense? Amia Srinivasin puts the dilemma this way:

Either effective altruism, like utilitarianism, demands that we do the most good possible, or it asks merely that we try to make things better. The first thought is genuinely radical, requiring us to overhaul our daily lives in ways unimaginable to most. . . . The second thought—that we try to make things better—is shared by every plausible moral system and every decent person.”¹³

I doubt there is a single answer that applies to all who consider themselves effective altruists. For one thing, EA is at least as much a movement as a theory, and as such it seeks to cast a wide net, embracing people with different underlying philosophies. (Consequentialists, in fact, can recommend such a move, on the grounds that it produces better consequences by not setting the bar too high or making the membership criteria too strict.) Remember MacAskill’s question: “Of all the ways in which we could make the world a better place, which will do the most good?”¹⁴ This formulation suggests that

¹¹ Ibid., 1; see also Unger, *Living High and Letting Die: Our Illusion of Innocence* (Oxford: Oxford University Press, 1996).

¹² Interestingly, in an earlier version of his essay McMahan writes: “my experience as a moral philosopher has been to find that common moral beliefs are often confused and inconsistent. . . . Whenever I consider a moral issue with care, I inevitably find the common sense view rather shallow. It is always possible to go deeper.” This passage might be thought to support my claim: the superficially common-sensical view expressed in Singer’s premise turns out on closer analysis to be confused and inconsistent. Of course, McMahan’s statement was meant to support the opposite conclusion: “I therefore think it is a mistake to suppose that the moral views of effective altruists can be rejected on the ground that they are more demanding than people now and in the past have thought that morality could be.” (Earlier version available from JL.)

¹³ Amia Srinivasin, “Stop the Robot Apocalypse,” *London Review of Books* 37, no. 18, September 24, 2015, at <http://www.lrb.co.uk/v37/n18/amia-srinivasan/stop-the-robot-apocalypse>.

¹⁴ MacAskill, *Doing Good Better*, 32.

making the world a better place is not mandatory, but that insofar as you adopt that goal you should make it as good as you can. But MacAskill also says that “Effective altruism consists of the honest and impartial attempt to work out what’s best for the world, and a commitment to do what’s best, whatever that turns out to be.”¹⁵ That sounds more like Singer’s demanding view.

Julia Wise, president of Giving What We Can and community liaison for the Centre for Effective Altruism, defends a more moderate position. She asserts that although cost-effectiveness is a useful tool that people should apply to solving global problems, it is not applicable to all parts of a person’s life:

not everything that you do is in the “effectiveness” bucket. I don’t even know what that would look like. . . . If I donate to my friend’s fundraiser for her sick uncle, I’m pursuing a goal. But it’s the goal of “support my friend and our friendship,” not my goal of “make the world as good as possible.” . . . That money is coming from my “personal satisfaction” budget, along with getting coffee with my friend. . . . I have another pot of money set aside for donating as effectively as I can. When I’m deciding what to do with that money, I turn on that bright light of cost-effectiveness and try [to] make as much progress as I can on the world’s problems.¹⁶

Is it acceptable, then, to have different pots of money—some to satisfy our personal desires, some to support relationships we care about, some for solving global problems like poverty and disease—with amounts allocated just as you please? Should Wise’s donation to the fundraiser for her friend’s sick uncle be understood simply in terms of “personal satisfaction,” not so different from getting coffee with a friend or buying a new dress? What about the pot of money “set aside for donating as effectively as I can”? Well, if the pot is for donating as effectively as you can, it’s tautological that with it you should donate as effectively as you can.

Presumably effective altruists do not mean to state a tautology, so let’s ask what **nontautological** recommendation might follow from their statements. One might be to be careful, when donating to causes like global poverty, to make sure the organization you support is not wasteful, corrupt, or inefficient, and that it is making good use of your donation. This is pretty uncontroversial, hardly enough to propel a movement meant to inspire people, as EA seems to be. Is it wrong to give to an organization that is not the most effective? Or, if not wrong, is it less morally good? How should a person decide

¹⁵ MacAskill, *Doing Good Better*, 11.

¹⁶ Julia Wise, “You Have More Than One Goal, and That’s Fine,” *Giving Gladly*, <http://www.givinggladly.com/>, February 19, 2019. Peter Singer writes about Wise in *The Most Good You Can Do*, 23–31, as does Larissa MacFarquhar in *Strangers Drowning: Grappling with Impossible Idealism, Drastic Choices and the Overpowering Urge to Help* (New York: Penguin Press, 2015), 71–87. MacFarquhar shows a different aspect of Wise: since childhood she had “believed that because each person was equally valuable she was not entitled to care more for her own well-being than for anyone else’s” (73). Perhaps she has made peace with these more radical and guilt-inducing views over the years. It’s also possible that, as an official representative of EA organizations, she realizes the need to encourage people to adopt its ways without inducing undue guilt.

how much to allocate to different causes, some of which are more effective than others? Effective altruists aim to persuade people to give much more than they currently do to organizations that effectively address global poverty (and other large problems). And this inevitably leads us to ask how demanding their view is.

3.

It's an interesting question why demandingness was not an issue for consequentialists until around the last third of the twentieth century—spurred no doubt in part at least by Singer's article. John Stuart Mill expressed a much more moderate view in *Utilitarianism*:

The multiplication of happiness is, according to the utilitarian ethics, the object of virtue: the occasions on which any person (except one in a thousand) has it in his power to do this on an extended scale—in other words, to be a public benefactor—are but exceptional; and on these occasions alone is he called on to consider public utility; in every other case, private utility, the interest or happiness of some few persons, is all he has to attend to.¹⁷

Henry Sidgwick took a similar line, asserting that “each person is for the most part, from limitation either of power or knowledge, not in a position to do much good to more than a very small number of persons.”¹⁸ The world Mill and Sidgwick inhabited was very different from ours, it's true; the efficacy of fairly ordinary people today may well be greater than it was in their day. But perhaps also they were self-deceived, or at least negligent in their thinking.

Mill softens the demandingness problem in his very definition of utilitarianism: “actions are right *in proportion* as they tend to promote happiness, wrong as they tend to produce the reverse of happiness.”¹⁹ Today this is often called scalar consequentialism: the more good you do the better; but it doesn't follow that you are morally *obligated* to do the most good.²⁰ Utilitarians can draw lines (on utilitarian grounds, of course) at places that evaluative terms like “wrong,” “impermissible,” “blameworthy,” and the like (including their positive counterparts) should mark.²¹

¹⁷ J. S. Mill, *Utilitarianism*, 2nd ed., edited by George Sher (Indianapolis: Hackett, 2001), 19.

¹⁸ Henry Sidgwick, *The Methods of Ethics* (New York: Dover, 1966) (republication of the 7th ed., 1907), 434. Sidgwick here seems not recognize the possibility of culpable ignorance.

¹⁹ Mill, *Utilitarianism*, 7. Emphasis added.

²⁰ Terms like “scalar consequentialism” are often technically defined. I am using the term in a broad sense, to include theories that draw no lines but just denote degrees (of goodness or badness), as well as those that do draw lines of the sort I describe.

²¹ Although I describe Mill here as a utilitarian, the case can be made that he is either a pluralistic consequentialist or not a consequentialist at all.

Sidgwick, by contrast, argues that “a Utilitarian must hold that it is always wrong for a man knowingly to do anything other than what he believes to be most conducive to Universal Happiness.”²² G. E. Moore thought similarly, writing, in *Principia Ethica*, that “the assertion ‘I am morally bound to perform this action’ is identical with the assertion ‘This action will produce the greatest amount of good in the Universe.’” Moore insists that the point is “demonstrably certain.”²³

The maximizing view became standard. In 1979, Brian Barry described it as “the time-bomb that has been ticking away ever since” and that “has at last blown up utilitarianism.”²⁴ But Barry was wrong; maximizing consequentialism is alive and well. For example, Shelly Kagan, a leading consequentialist, takes consequentialism to mean that “Agents are morally required to perform the act that will lead to the best results overall.”²⁵

Neither Singer nor MacAskill uses the term “duty” or “obligation” in these works, nor do they say you are wrong if you do not do what’s best; rather, they say you *ought* to do the best thing you can. “Ought” is an ambiguous word. Perhaps, then, they are not maximizing consequentialists but instead accept Mill’s scalar view. If so, doing less than the best is not necessarily wrong. Nevertheless, Singer asks us to do the most good we can, and other effective altruists suggest similar things. In at least one place Singer implies that one who does not aim to bring about the very best consequences is morally deficient and open to criticism.²⁶ So there is a crucial ambiguity running through the pronouncements of effective altruists.

But unless it is so defined, there is nothing in consequentialism that entails the maximizing view. “Consequentialism . . . is simply the view that normative properties depend only on consequences.”²⁷ “A moral theory is a form of consequentialism if and only if it assesses acts and/or character traits, practices, and institutions solely in terms of the

²² Sidgwick, *Methods of Ethics*, 494.

²³ G. E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1971; originally published 1903), 147. Should there be any doubt about his meaning, a paragraph later he emphasizes that “Our duty . . . can only be defined as that action, which will cause more good to exist in the Universe than any possible alternative.”

²⁴ Brian Barry, “And Who Is My Neighbor?,” *Yale Law Journal* 88 (1979), 639n37 (a review essay of Charles Fried, *Right and Wrong*). I discuss this matter further in Judith Lichtenberg, “The Right, the All Right, and the Good,” *Yale Law Journal* 92 (1983) (a review essay of Samuel Scheffler, *The Rejection of Consequentialism*). Barry tells us that the term “consequentialism” first appeared in G. E. M. Anscombe’s “Modern Moral Philosophy,” *Philosophy* 33 (1958); Anscombe takes credit for it on p. 10 of the essay.

²⁵ Shelly Kagan, *The Limits of Morality* (Oxford: Clarendon Press, 1989), xi.

²⁶ Singer asserts that, unlike effective altruists who donate to one or two charities about which they have evidence of effectiveness, “those who give small amounts to many charities are not so interested in whether what they are doing helps others—psychologists call them warm glow givers” (*The Most Good You Can Do*, 5). So on this view those who spread their donations more broadly are not simply less effective; they act out of self-interest. Singer gives no warrant for maligning their motives. At the very least one would have to show that such givers are culpably ignorant—that they have been informed or ought to know about the “scientific findings” of EA and have ignored its message.

²⁷ Walter Sinnott-Armstrong, “Consequentialism,” *Stanford Encyclopedia of Philosophy*, at <https://plato.stanford.edu/archives/sum2019/entries/consequentialism/>

goodness of the consequences.”²⁸ Consequentialism entails that the more good you do the better, but it is not part of its meaning that you are morally required to do what is best. And I find no convincing arguments by consequentialists for the conclusion that maximizing the good is morally required. (As I will suggest later below, it’s hard to know what would count as a persuasive argument for it.) Where we should draw the lines between acts that are prohibited, permissible, praiseworthy, and the like will depend, according to consequentialism, on what promotes the best consequences.

Nonmaximizing consequentialism is defanged of a property that makes consequentialism highly controversial. Still, essential to consequentialism is the idea that it is always morally *preferable* to bring about better consequences than worse ones.

This view has often been challenged on deontological grounds, by those who believe it is sometimes *wrong* to bring about more good rather than less. But there are other reasons to doubt it too.

4.

Consider Jane, a chemistry professor in her forties with a successful career and a good salary in a respected university. She gives generously to causes devoted to alleviating global poverty. She does not impoverish herself in the process, but she donates much more than the average American household and acts in line with recommendations of the EA movement, carefully researching nongovernmental organizations (or relying on others, like GiveWell, to do so) and giving her money to organizations with a proven track record of success in alleviating the worst problems facing the global poor.

But over the years Jane has developed interests in other major social problems too. In particular, she has become deeply disturbed by the glaring social and economic inequalities in the United States and the legacy of racism and discrimination underpinning them. So disturbed is she, in fact, that she eventually decides to quit her academic job to work with at-risk teenagers in her city, mostly members of minorities—working to keep them in school, off drugs, out of trouble, and on track to go to college or learn a trade. She also works part-time for a nonprofit organization with similar goals. Together, these activities amount to a full-time job and then some. But she makes much less than what she made as an academic, significantly diminishing the resources she has to give to global

²⁸ Brad Hooker, “Rule Consequentialism,” *Stanford Encyclopedia of Philosophy*, at <https://plato.stanford.edu/archives/win2016/entries/consequentialism-rule/>. Some would describe this view as global consequentialism (see Greaves, Chapter 22, this volume) rather than consequentialism full-stop. Douglas Portmore describes himself as an act consequentialist; he believes we must judge acts in consequentialist terms but denies that whether a person is blameworthy or virtuous or whether a belief is rational depends on whether blaming the person or having certain character traits or forming the belief would maximize the good. See Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford: Oxford University Press, 2011). Here I stick to the more standard and old-fashioned understandings of consequentialism.

poverty relief. Over the twenty or more years she would likely have continued as an academic, those foregone contributions, according to MacAskill's calculations, could have saved some significant number of lives. Jane's work saves no lives (that we know of). And let's suppose that despite the good it achieves it does not reduce human suffering as much as those donations would have (however we choose to measure these things).

Many would consider Jane admirable, but by the standards of EA she is open to criticism. Consider what the organization 80,000 Hours advises to those who wish to relieve human suffering. First, figure out which problems are "large in scale, solvable and neglected." As of this writing, the organization lists "positively shaping the development of artificial intelligence," "biorisk reduction," nuclear security, and climate change.²⁹ Health in poor countries is also a priority, although it seems to have fallen a few notches in the last several years, perhaps because it is less neglected as a field than some of the others on the list.³⁰ Effective altruists also note that global health and poverty alleviation will always take priority over their domestic analogues, largely because dollars going to the least well-off produce more utility than dollars going to those higher on the socioeconomic scale. The poorest 5 percent of Americans are "at the 68th percentile of the world income distribution."³¹

80,000 Hours' next recommendation is to engage in research, or government or policy work, in one of those needed areas; or to work at an effective nonprofit; or to "apply an unusual strength to a needed niche."³² Another strategy is "earning to give"—taking a high-paying job, then donating a good chunk of one's salary to an urgent problem. In his book MacAskill describes Greg Lewis, an idealistic doctor who decided not to practice medicine in a poor country but instead to become a medical oncologist in the United Kingdom so that he could donate half his \$200,000 earnings to global poverty relief.³³ Singer begins his book with a discussion of Matt Wage, a brilliant former student of his at Princeton who declined to continue in philosophy despite an offer from Oxford and instead went to work for an arbitrage trading firm on Wall Street in order to be able to donate six-figure sums to highly effective charities.³⁴

Many criticisms can and have been brought against EA, attempting to show its methods are not "better." That it is elitist, concentrating almost exclusively on how highly educated people can make a difference (the movement's home is in Oxford, and it shows); that it is individualistic and apolitical, whereas the problems it addresses are

²⁹ <https://80000hours.org/key-ideas/>

³⁰ Other things being equal, you can make more of a difference when an important issue is neglected than if it isn't. Here I mostly focus on global health and poverty, which have been and remain centerpieces of the EA movement.

³¹ Branco Milanovic, *The Haves and the Have-Not: A Brief and Idiosyncratic History of Global Inequality* (New York: Basic Books, 2011), 116.

³² <https://80000hours.org/key-ideas/>

³³ MacAskill, *Doing Good Better*, 76–77.

³⁴ Singer, *The Most Good You Can Do*, 3–4. Some will worry about the harms such careers might cause. In his book MacAskill describes a documentary filmmaker who criticizes one of his own subjects, a cosmetic surgeon to the stars, for wasting his talent rather than saving lives. MacAskill argues that the filmmaker's attitude "is misplaced. It's the cosmetic surgeon's decision about how to spend his money that really matters" (*Doing Good Better*, 78). (According to Srinivasan, MacAskill no longer recommends that people take jobs that cause direct harm. "Stop the Robot Apocalypse.")

political and must be solved collectively by governments and other groups; that poverty is best addressed by economic development; that the techniques of EA are paternalistic and undemocratic, treating those it helps as passive and irrelevant.³⁵ To such criticisms Singer responds: “Effective altruism cannot be refuted by evidence that some other strategy will be more effective than the one effective altruists are using, because effective altruists will then adopt that strategy.”³⁶ If political action is a better means to improving welfare than individual giving, that’s what we should promote. If private investment by Western conglomerates works best, do that. But there are, it seems, no intrinsic advantages of these approaches.

This response can make it difficult for criticism of EA to stick. Is it an adequate answer? Let me explain why I believe it is not.

5.

One reason has to do with how much we can pack into the conception of the good without trivializing consequentialism. If we reject the simpler utilitarian conception—the good is pleasure, happiness, or well-being; the bad is pain, unhappiness, or ill-being—and add a bunch of other goods such as autonomy, dignity, integrity, self-determination, fidelity, democracy, and human rights, then we will need some way of weighing the various values against each other in cases where they conflict, as they inevitably will some of the time, to decide which takes priority in what circumstances. For consequentialism to be a distinctive moral theory there must ultimately be a single scale along which to measure these different values. Otherwise it resembles an intuitionism that judges (non-scientifically or nonquantitatively) which value takes precedence when.³⁷

Effective altruists and consequentialists seem highly optimistic about the prospects for measuring, quantifying, and comparing these values—including the less hedonistic ones—to produce the right answer about what to do. Of course, effective altruists are rational, scientific-minded people who embrace fallibilism; even in the few years since the movement started they have changed their minds quite a bit about what strategies to pursue. Nonetheless, they seem to engage in predictions with great confidence. For example, among GiveWell’s most recommended charities (there are only eight as of this writing) are Evidence Action’s Deworm the World Initiative and the Schistosomiasis Control Initiative, both of which seek to eradicate intestinal worms, which afflict more

³⁵ For such criticisms see, e.g., the brief essays in the symposium on EA, “The Logic of Effective Altruism” (*Boston Review* 40, July/August 2015), by Daron Acemoglu, Emma Saunders-Hastings, Angus Deaton, Jennifer Rubenstein, Leila Janah, Larissa MacFarquhar, and others.

³⁶ Peter Singer, “The Logic of EA,” *Boston Review* 40, July/August 2015, 31.

³⁷ Douglas Portmore points out that “Even hedonists have to rely on intuitions. For instance, even the simplest Benthamite hedonist must rely on intuition to determine how to make a trade-off between the duration of pleasure and the intensity of pleasure” (personal correspondence). Some will take this as a defense of pluralistic consequentialism, others as a rebuttal of even simple utilitarianism.

than two billion people in poor countries.³⁸ But in 2015, Cochrane, an organization that evaluates health and medical research data, cast doubt on the efficacy of mass deworming programs; its findings were confirmed in a report published in the *Lancet* in 2017.³⁹ Angus Deaton, a Nobel Prize-winning economist who is critical of EA, attributes the movement's overconfidence in this and other cases to overreliance on randomized experiments, which "consider only the immediate effects of the interventions, not the contexts in which they are set. Nor, most importantly, can they say anything about the wide-ranging unintended consequences."⁴⁰

James Lenman describes the problem decision makers face as "massive and inscrutable causal ramification."⁴¹ Thus, he argues, the kinds of predictions on which consequentialists rely and judge others' behavior are highly questionable. Consequentialists may respond that this is not a problem faced by them alone: everyone, whatever their theories, must make decisions under uncertainty. But consequentialists face the problem in a more severe form, since for them consequences are the *only* things that matter and *all* the consequences matter. Nothing else counts. GiveWell, for example, whether it explicitly embraces consequentialism or not, is vulnerable to these charges partly because it generally recommends that people put all their charitable eggs in just a few baskets. Less confidence might lead one to spread one's donations more widely.

6.

But let's leave this difficulty aside. The main problem I want to address arises from the fact that even a nonmaximizing, scalar consequentialist must be committed to the view that bringing about more good, whatever that good is, is always morally better than bringing about less. That is inherent in consequentialism. So, for example, what Jane does is worse than if she had gone to work for one of the organizations 80,000 Hours recommends or if she had "earned to give" (assuming that either or both of these would have produced better overall results). Her chosen course may be better than doing nothing, but it is far less morally good than other options she could have taken.

³⁸ GiveWell, "Top Charities," at <https://www.givewell.org/charities/top-charities>

³⁹ See Cochrane Review, "Deworming Children in Developing Countries," https://www.cochrane.org/CD000371/INFECTN_deworming-school-children-developing-countries; Jason R. Andrews et al., "The Benefits of Mass Deworming on Health Outcomes: New Evidence Synthesis, the Debate Persists," *Lancet* 5, no. 1, January 1, 2017, [https://www.thelancet.com/journals/langlo/article/PIIS2214-109X\(16\)30333-3/fulltext](https://www.thelancet.com/journals/langlo/article/PIIS2214-109X(16)30333-3/fulltext); Susan Brink, "The Debate Is On: To Deworm or Not to Deworm," NPR, November 15, 2015, https://www.npr.org/sections/goatsandsoda/2015/11/15/452298843/the-debate-is-on-to-deworm-or-not-to-deworm?utm_source=facebook.com&utm_medium=social&utm_campaign=npr&utm_term=nprnews&utm_content=2041?utm_source=facebook.com&utm_medium=social&utm_campaign=npr&utm_term=nprnews&utm_content=2041

⁴⁰ Deaton, "The Logic of EA," *Boston Review* 40, July/August 2015, 21.

⁴¹ James Lenman, "Consequentialism and Cluelessness," *Philosophy & Public Affairs* 29 (2000), 348.

The question is whether this view, which I believe is at odds with our common-sense beliefs, can be made convincing by means other than building it into the definition of consequentialism. From a consequentialist point of view, what Greg and Matt do just *is* morally better than what Jane does. Here's another story that MacAskill tells. In 2009, while germinating the concept of EA, he visited a hospital in Ethiopia that treats obstetric fistulas, a condition resulting from childbirth in young and malnourished women that causes "permanent incontinence of urine and/or feces" and pariah status. According to the Fistula Foundation, which funds the hospital, "A majority of women who develop fistulas are abandoned by their husbands and ostracized by their communities because of their foul smell." At the hospital MacAskill met some of the women who suffered from this condition. But several years later he concluded that although the organization was repairing fistulas at low cost and saving these young women from terrible fates, others working on different issues were making a bigger impact and that they should get his donations instead. (The methods used in determining impact are described in the book in detail.) MacAskill explains that by donating to the Fistula Foundation instead of a different organization he thought he "would be privileging the needs of some people over others merely because I happened to know them," and that it "was arbitrary that I'd seen this problem close up rather than any of the other problems in the world."⁴²

Effective altruists often applaud MacAskill's approach, admiring the cool rationality that considers personal attachments arbitrary. Singer tells us that "many of the most prominent effective altruists have backgrounds in or are particularly strong in areas that require abstract reasoning."⁴³

Effective altruists are right that people are often led astray, in a variety of ways, by their emotions and personal attachments, and that these can lead to pernicious biases. But critics may nevertheless find MacAskill's approach chilling. As Larissa MacFarquhar puts it, effective altruists fail to understand "that, to many people, to suppress emotional connection to make way for a more rational altruism is to crush their moral roots."⁴⁴

There are several different worries implicit here. Philosophers have long criticized consequentialism on the grounds that it ignores the moral value of partiality. Even on the scalar version, it is always *permissible*, never wrong, to bring about more good rather than less, and that can mean neglecting those near and dear in favor of strangers. But, as Diane Jeske argues elsewhere in this volume, "many of us cannot help but think that choosing to benefit our friend [rather than a stranger] is not only morally permissible but, in fact,

⁴² MacAskill, *Doing Good Better*, 41–42. See also Wise, "No One Is a Statistic," *Giving Gladly*, October 10, 2018. Wise describes "a mother of seven who lives in rural Guatemala and has cervical cancer. The doctor treating her knows that screening other women for cancer is more cost-effective than treating this woman, and that the community doesn't have enough money to fully fund both." These other women might be considered "mere statistics," and treating the one who already has cancer might seem more humane. Wise wisely expresses the agonizing dilemma: "Here's the thing about those 'statistics': they're all individuals."

⁴³ Singer, *The Most Good You Can Do*, 89.

⁴⁴ MacFarquhar, "The Logic of Effective Altruism," 25. For an account of altruists excruciatingly in touch with their moral roots, see MacFarquhar, *Strangers Drowning*.

morally required.”⁴⁵ Jeske, like many others, defends “partialism,” according to which it is “(not merely psychologically understandable but) morally correct to favour one’s own . . . [i.e.] those to whom the agent has some special relationship or personal tie.”⁴⁶

Two features of the kinds of examples offered by partialists are important. First, they involve cases in which one confronts the choice to benefit (or prevent harm to) an intimate rather than a stranger. Second, the partialist casts the matter in deontic terms: we have a *duty* to favor those with whom we have a relationship; it would be wrong not to.⁴⁷

Jane’s case exhibits neither of these features. Jane has no prior relationship with the children she works with; she makes a decision *ex ante* to work with them rather than to “earn to give.” And, partly for that reason, it would be odd to say she is required to work with them rather than do something more utility-producing. In these cases, deontic language does not describe our attitudes. It would be more natural to say that Jane, or what she does, is *morally admirable*. We think highly of the traits associated with her behavior. We esteem certain human characters and characteristics. And we may believe that Jane’s behavior is *no less admirable than earning to give*.

Can effective altruists and consequentialists make sense of these beliefs? They may resist what seems to them like a naïve conflation of “does more good” with “is more admirable” or “is a good person.” But consequentialists must interpret the judgment of human character, like everything else, in consequentialist terms.⁴⁸ The strategies they may use to incorporate the common-sense view that Jane is no less admirable than (what I shall call) the philanthropist resemble those employed against deontologically minded critics. Just as consequentialists can argue that in the long run the world will be a better place if people take care of their own and are partial to those near and dear to them, they may say that the world will go better if some people choose local, hands-on efforts such as working with at-risk teenagers or illiterate adults, addressing prison reform or homelessness in the United States, reporting on gang violence, or working for the rights of transgender people. The character traits possessed by such people are admirable because they tend to produce good consequences.

Yet today the claim that from a consequentialist point of view local efforts might produce just as much good as global efforts seems implausible. (And, indeed, effective altruists explicitly deny it.) If we compare the numbers of people who suffer globally from poverty and disease to, say, the numbers of Americans who do, and the extent to which

⁴⁵ Jeske, Chapter 12, this volume.

⁴⁶ John Cottingham, “Partialism, Favouritism, and Morality,” *Philosophical Quarterly* 36 (1986), 357–358; quoted in Jeske, Chapter 12, this volume.

⁴⁷ Even the weaker claim that one is permitted to be partial is expressed in deontic terms. But see Bernard Williams for the “one thought too many” argument (“Persons, Character, and Morality,” in Williams, *Moral Luck* [Cambridge: Cambridge University Press, 1981], 18). Williams famously asserts that if a person had to choose whether to save his wife or a stranger, it would be “one thought too many” for him to think “It is morally permissible for me to save my wife.” Unlike the standard partialist critique and the one I make here, in such situations Williams objects to thinking in moral terms at all.

⁴⁸ As I suggested in note 28, some new and sophisticated versions of consequentialism try to avoid this.

these problems could be remedied, the former will swamp the latter by a good margin.⁴⁹ Exactly so!, the effective altruists may say. This shows, they think, that those who earn to give *are* more admirable than those who work hands-on.

To avoid this conclusion, consequentialists might argue that the traits that lead ordinary people to be generous or self-sacrificing or to help people in their own communities represent traits that, over the long haul of human history, have produced highly beneficial consequences. We should not be short-sighted; the benefits of encouraging the local, hands-on approach, and of doing what you can with what you have, are greater than the comparison of the contemporary global and local poor would suggest.

7.

How long a time frame must a consequentialist consider, then, in recommending courses of action? This draws us near to the other approach consequentialists have taken to soften the clash between their views and the common-sense commitment to partiality and localism: to adopt some form of indirect consequentialism, such as rule consequentialism or motive consequentialism. Rule consequentialism, the variation that has been most extensively developed, says we should adopt those general rules adherence to which would maximize the best consequences. What is the life expectancy of a rule? Must we change the rules from those that may have worked for most of human history? In its pure form, rule consequentialism is a very different beast from act consequentialism that can lead to wholly different recommendations for action. It has a Kantian flavor, asking “What if everybody did that?” Despite its appeal, rule consequentialism is often thought to be, as Philippa Foot puts it, an “unstable compromise” whose very consequentialism can be called into question. One threat is that it “collapses” into act consequentialism.⁵⁰ In any case, I find no evidence that contemporary effective altruists accept indirect consequentialism.⁵¹

⁴⁹ I leave aside here the problem of comparing absolute and relative poverty and its import: whether even if, e.g., poor Americans have more to eat, their deprivations relative to others in their society is a further element that must be considered in evaluating their absolute level of well-being. See Lichtenberg, *Distant Strangers*, chaps. 5 and 6.

⁵⁰ Philippa Foot, “Utilitarianism and the Virtues,” Presidential Address to the 57th Annual Meeting of the Pacific Division of the American Philosophical Association (*Proceedings and Addresses of the APA* 57 [1983], 273. For the collapse argument see David Lyons, *Forms and Limits of Utilitarianism* (Oxford: Oxford University Press, 1967); and the discussion in Hooker, Chapter 23, this volume.

⁵¹ There are other important distinctions bearing on these questions that I cannot pursue here. One is between consequentialism as a criterion of rightness versus consequentialism as a decision procedure (see Peter Railton, “Alienation, Consequentialism, and the Demands of Morality,” *Philosophy & Public Affairs* 13 [1984]). Another is between how I should go about deciding what I myself should do versus what acts or procedures I recommend to others (a significant distinction especially for writers, teachers, and public intellectuals). Related is the distinction between a public and an esoteric morality (“a Utilitarian may reasonably desire, on Utilitarian principles, that some of his conclusions should be rejected by mankind generally”; Sidgwick, *Methods of Ethics*, 490).

Another strategy consequentialists might adopt is to include traits like admirability and goodness as intrinsic goods to be included among all the other things we think are good. It might not be any more difficult, or any more necessary, for them to explain on what grounds we include such traits than it would be for anyone else, whatever their moral theory. But this will not solve the problem, because these goods will almost certainly be swamped by others (full stomachs, long and healthy lives). Here again effective altruists and other consequentialists may embrace the consequentialist implication—what we might call the QED response—but others will remain unsatisfied.

One unfortunate upshot of the effective altruist's preference for big philanthropy and highly skilled careers is to downgrade the lives and works of those without the skills to attain such careers and unable to earn large sums of money to donate. They might give generously, given their means, but their means severely limit their impact. Can consequentialists avoid the embarrassing conclusion that such people are less admirable than rich people who have a much greater impact?

We can assume they will not want to say that. Like the rest of us, they are probably drawn to the intuitive idea that one's "goodness" (admirability, virtue, decency, moral compass) is largely a function of what one does with what one's got, how much one does for others, and so on. But it is not clear how to incorporate any of these convictions on consequentialist grounds.

8.

It's not news that pure consequentialism conflicts with various of our common-sense beliefs. Its clash with deontological intuitions has dominated discussions of moral theory for decades. Here I have focused on a less well-trodden conflict: with our judgments about what constitutes moral virtue and admirability.

Effective altruists and other consequentialists are right to insist that the mere fact that our "intuitions" tell us that people like Jane are as good as the generous capitalist is no argument at all. If common sense were the last word, then slavery would have been justified to those who thought its legitimacy was intuitively clear.

But common-sense beliefs about morality are not always suspect. Common sense is especially untrustworthy when what it endorses happens to coincide with our interests. When it justified slavery to slaveholders, there were deep reasons to distrust it. Moreover, there were many reasons for thinking slavery was immoral. But no self-interested motive is apparent when common sense tells us that we should not knowingly convict the innocent or that civil rights workers are admirable. In such cases we have much less reason to distrust our intuitions.

Of course, our self-interest *is* at stake when we reject a theory that demands or even just recommends that we do more than would be convenient for us. Yet many moral and political theories, outlooks, and religions express ideals, aspirations, or even moral requirements that far surpass what most people do. What makes consequentialism

unacceptable is not this, but rather its commitment to calculating and quantifying the Good, which guarantees that many actions, traits, and habits that we value will be swamped by the numbers.

Consequentialists who admit the intrinsic value of traits like Jane's can recognize this outcome as an unfortunate but inevitable by-product of their theory. But some think it is enough to discredit any view that takes how the numbers fall out as decisive.

This disagreement cannot be rationally decided.

9.

Is EA necessarily consequentialist? Clearly, there are different versions of EA. No doubt not every effective altruist is a consequentialist—for one thing, some may never have heard of consequentialism or may not have a formally coherent moral outlook; some may simply think that those who can ought to do more for those in need. But of those who have thought in these terms I expect (and have argued here) that most are consequentialists, not necessarily in the maximizing sense but in the more modest sense of believing that producing more good is always morally better than producing less. And that in itself is controversial.

But EA is both a *movement* and a *theory* or *view*. Insofar as EA is a movement, its members will want to get other people to join. They may well decide that the best way to do that is by emphasizing aspects of the movement that are appealing, satisfying, and not too onerous—that do not much disrupt people's existing aims, values, and habits. Stern moral imperatives that challenge people's moral decency may be counterproductive. Effective altruists may well be committed to a theory, a set of moral truths—perhaps even that one ought to do as much as one possibly can to increase well-being overall—but may recognize that expressing the theory widely and publicly will not necessarily help to bring about its goals.⁵² And effective altruists want, after all, to be effective.⁵³

⁵² See Sidgwick, *Methods of Ethics*. Sidgwick recognizes that such conclusions are “of a paradoxical character” and that “the moral consciousness of a plain man broadly repudiates the general notion of an esoteric morality” (489–490).

⁵³ This paper draws in a few places from Judith Lichtenberg, “Peter Singer's Extremely Altruistic Heirs,” *The New Republic*, November 30, 2015, and Judith Lichtenberg, “The Right, the All Right, and the Good.” I am grateful to Douglas Portmore and Victor Kumar for helpful comments on an earlier draft.